UNICAMP

Universidade Estadual de Campinas

Faculdade de Ciências Aplicadas

Maria Paula Buelvas Padilla

A study of flexibility and the chaining principle in the mid-term nurse scheduling problem

Um estudo de flexibilidade e o princípio da regra da cadeia para o problema de escalonamento de enfermeiras

Limeira

2021

UNICAMP

Universidade Estadual de Campinas

Faculdade de Ciências Aplicadas

Maria Paula Buelvas Padilla

A study of flexibility and the chaining principle in the mid-term nurse scheduling problem

Um estudo de flexibilidade e o princípio da regra da cadeia para o problema de escalonamento de enfermeiras

Dissertação apresentada à Faculdade de Ciências Aplicadas como parte dos requisitos exigidos para a obtenção do título de Mestra em Engenharia de Produção e de Manufatura na área de Pesquisa Operacional e Gestão de Processos.

Orientador: Ph.D. Diego Jacinto Fiorotto

Este exemplar corresponde à versão final da dissertação defendida pela aluna Maria Paula Buelvas Padilla, e orientada pelo Prof. Dr. Ph.D. Diego Jacinto Fiorotto.

Limeira

2021

Ficha catalográfica
Universidade Estadual de Campinas
Biblioteca da Faculdade de Ciências Aplicadas
Renata Eleuterio da Silva - CRB 8/9281

B862s

Buelvas Padilla, Maria Paula, 1995-
A study of flexibility and the chaining principle in the mid-term nurse scheduling problem / Maria Paula Buelvas Padilla. – Limeira, SP : [s.n.], 2021.

Orientador: Diego Jacinto Fiorotto.
Dissertação (mestrado) – Universidade Estadual de Campinas, Faculdade de Ciências Aplicadas.

1. Escalonamento. 2. Enfermagem. 3. Treinamento. 4. Programação inteira. I. Fiorotto, Diego Jacinto, 1987-. II. Universidade Estadual de Campinas. Faculdade de Ciências Aplicadas. III. Título.

Informações para Biblioteca Digital

Título em outro idioma: Um estudo de flexibilidade e o princípio da regra da cadeia para o problema de escalonamento de enfermeiras
Palavras-chave em inglês:
Scheduling
Nursing
Training
Integer programming
Área de concentração: Pesquisa Operacional e Gestão de Processos
Titulação: Mestra em Engenharia de Produção e de Manufatura
Banca examinadora:
Diego Jacinto Fiorotto [Orientador]
Karim Yaneth Pérez Martínez
Carla Taviane Lucke da Silva Ghidini
Data de defesa: 30-09-2021
Programa de Pós-Graduação: Engenharia de Produção e de Manufatura

Identificação e informações acadêmicas do(a) aluno(a)
- ORCID do autor: https://orcid.org/0000-0003-2184-3448
- Currículo Lattes do autor: http://lattes.cnpq.br/0609928515109252

UNICAMP

Universidade Estadual de Campinas

Faculdade de Ciências Aplicadas

Maria Paula Buelvas Padilla

A study of flexibility and the chaining principle in the mid-term nurse scheduling problem

Um estudo de flexibilidade e o princípio da regra da cadeia para o problema de escalonamento de enfermeiras

Dissertação apresentada a Faculdade de Ciências Aplicadas da Universidade Estadual de Campinas como parte dos requisitos exigidos para a obtenção do título de Mestra em Engenharia de Produção e Manufatura na área de Pesquisa Operacional e Gestão de Processos[1].

Banca Examinadora

Prof. Dr. Diego Jacinto Fiorotto
UNICAMP - Limeira

Prof. Dra. Karim Yaneth Pérez Martínez
HEC - Montréal

Prof. Dra. Carla Taviane Lucke da Silva Ghidini
UNICAMP - Limeira

Limeira-SP
30 de setembro 2021

[1] A Ata da defesa com as respectivas assinaturas dos membros encontra-se no SIGA/Sistema de Fluxo de Dissertação/Tese e na Secretaria do Programa da Unidade.

Acknowledgements

The work you are about to read was written at a difficult time for me and, probably, most people. I started writing it in a moment when I wasn't sure if and when I would come back home and hug the people I love the most. Going through the entire roller coaster of emotions and finishing this work would not have been possible without the support and love of the wonderful people that God placed in my path. I would like to thank them in the following paragraphs.

I would like to express my gratitude to my advisor Diego Jacinto Fiorotto for presenting me the opportunity to be part of the *Engenharia de produção e Manufatura* program at UNICAMP, for his ideas, advice, and suggestions that allowed the completion of this work. My immense gratitude also goes to my hero without a cape, Gabriel de Souza Amaro. *Meu amigo!* this work would not have been possible without your help, guidance, and patience. Thank you for talking with me about life, research, encouraging my ideas, guiding me to solve my doubts and problems, and all that valuable time dedicated to helping me correct my messy code, and practice my portuguese.

Special thanks to the Espahnolito team Alejandra, Jhonatan, and Luis Fernando for their friendship, good company, and making me feel at home when I missed it the most. Also to the girls of the *República flor de Lis*, Carol, Juliana for receiving me with open arms and to Mayara and Mauricio for always being there to help me out, share experiences, and take the risk of visiting me in the middle of such a difficult time. Thanks to my friends in Colombia and Brasil Angie (Ma Bestie), Jesús, Oscar, Carlos, and Giannini for remembering me in the distance. Thank you Lu, for sharing your days with me and letting me share mine with you.

I would like to thank my teachers Helman Enrique Hernandez Riaño and Jorge Mario López Pereira for motivating me to learn a little bit more about research during my graduation, and help me out sharpening my analytical skills.

Finally, I am eternally grateful to my parents Sergio and María for always supporting me with their unconditional love and understanding, to my sister Vanessa for frequently knocking on my door with her music and a dose of joy to brighten my life, and to my grandmother for her love and patience. Thanks to all my family for being there for me all this time.

Special thanks to the Fundação de Desenvolvimento da Unicamp - FUNCAMP for the financial support to carry out this research through agreement no. 519.292-1, Auxílio Sol. 3082/19.

A special mention to Oliver*, o gato mais chato da Av. Cônego Manoel Alves, for the company.*

A todos, ¡infinitas gracias!

Resumo

O escalonamento de pessoas é um problema comum que as instituições de saúde freqüentemente enfrentam em todo o mundo. O problema consiste na designação de enfermeiras para trabalhar em turnos considerando de uma série de restrições relativas às habilidades da força de trabalho, preferências dos funcionários, tempo, normas legais, entre outras. O contexto atual da saúde, onde se espera um aumento da demanda e, ao mesmo tempo, uma redução da força de trabalho, juntamente com o efeito que o escalonamento tem sobre aspectos como custos, aposentadorias antecipadas, rotatividade e satisfação no trabalho em geral, tornou-se uma grande motivação para estudar este problema. O *cross-training*, entendido como o treinamento de pessoal para trabalhar em outras unidades diferentes de suas unidades dedicadas, tem sido estudado em diferentes cenários de produção e serviços, incluindo a saúde. O principal objetivo deste trabalho consiste em estender os trabalhos existentes na literatura que abordam a flexibilidade através do *cross-training* no problema de escalonamento de enfermeiras. Estudamos as formas existentes para incluir este tipo de flexibilidade e propomos duas extensões de um modelo e uma heurística para incorporá-lo. Uma série de experimentos computacionais sobre diferentes problemas mostram que as extensões propostas são funcionais e podem produzir soluções para o problema de escalonamento de enfermeiras abordado. A partir dos experimentos executados obtivemos alguns *insights* interessantes sobre a flexibilidade mediante cross-training. Em primeiro lugar, nossos resultados reafirmam as conclusões da literatura de que é mais vantajoso investir em intensidade, que é o número de enfermeiros treinados para trabalhar en outas unidades, do que investir em *breadth*, ou seja, o número de unidades adicionais para as quais uma enfermeira é treinada. Também constatamos que é importante considerar a ordem das unidades na utilização da regra da cadeia, o que leva a um aumento dos benefícios do cross-training, principalmente quando o investimento disponível para novos treinamentos é relativamente pequeno. Finalmente, com diferentes níveis de investimento foi observado um trade-off entre flexibilidade e proporção de tempo de trabalho dedicado para cada unidade em alguns problemas, algo que pode ser relevante em casos em que a eficiência dos enfermeiros é reduzida enquanto trabalham em unidades secundárias, porque, como a teoria sugere, isto pode ter um impacto negativo na qualidade do atendimento.

Palavras-chaves: Escalonamento; enfermagem; treinamento; programação inteira

Abstract

The rostering or scheduling of personnel is a common problem that health care institutions frequently face worldwide. It consists of assigning nurses to shifts to satisfy a series of constraints relative to the workforce's skills, employee preferences, time, legal regulations among others. The present context of healthcare, where it is expected an increase for demand and at the same time a reduction of the workforce, together with the effect that scheduling has over aspects like costs, early retirements, turnovers, and job satisfaction in general, has become a great motivation for studying this problem. Cross-training, which consists of training nurses to cover additional units, has been studied in different sectors from manufacturing to healthcare as a flexibility method that can bring good benefits. The main objective of this work is to extend existing works in the literature that address flexibility through cross-training in the mid-term nurse scheduling problem. We study existing forms to include this type of flexibility and propose two model extensions and a heuristic to incorporate it. A series of computational experiments over different problems show that the proposed extensions are functional and can produce solutions for the mid-term nurse scheduling problem. From the executed experiments, we derive some interesting insights about cross-training flexibility. First, our results reaffirm findings in the literature that it is more beneficial to invest in intensity, which is the number of cross-trained nurses, than investing in breadth i.e. the number of additional units a nurse is trained to. We also found that it is important to consider the order of units while using the cross-training policy known as chaining as additional benefits can be grasped specially when the available investment for cross-training is relatively low. Finally, with different levels of investment a trade-off between flexibility and the proportion of dedicated working time for every unit in some problems was observed, something that can be relevant in cases when the efficiency of nurses is reduced while working on secondary units, because, as the theory suggests, this can have a negative impact in the quality of care.

Keywords: Scheduling; nursing; training; chaining; integer programming

List of Figures

List of Tables

Contents

CHAPTER 1

Preliminars

In this work, we study flexibility in the mid-term nurse scheduling problem. We first state our motivation to study this problem and later organize the relevant literature on flexibility in various sectors. Then, we develop extensions for an existing model in the literature and perform experiments to analyze some well-known cross-training indicators and policies to derive insights. We briefly describe the content of each chapter as follows.

Chapter 2 starts by stating our motivation to study flexibility in healthcare, particularly in nurse scheduling, then introduces basic concepts to ease the comprehension of this research by defining and representing basic terms like flexibility, cross-training (CT), intensity, breadth, chaining, among others. In Chapter 3, the current status of the literature is presented. We provide a general definition of the nurse scheduling problem and show the reader relevant works by authors who have studied the problem over the years. Then, we focus on the study of flexibility in various sectors, as they serve as a general framework to study flexibility, to later center on flexibility in healthcare, describing how flexibility is currently managed in this sector, and how authors have addressed it. We finish this chapter by commenting on cross-training in the nurse scheduling problem and stating this work's contribution to the literature.

Chapter 4 introduces the formulations used throughout this work. We consider a base model in the recent literature and then propose variations that we later use to study cross-training flexibility in the nurse scheduling problem. In Chapter 5, we propose, perform and discuss a series of numerical experiments using the previous models. Finally, Chapter 6 summarizes the findings of our work and proposes future research directions.

CHAPTER 2

Introduction

2.1 Motivations to study flexibility in nurse scheduling problems

As the most representative and critical workforce in the healthcare sector, nurses will become highly demanded in the following years as the population ages and new chronic and degenerative diseases appear, increasing the demand for more health care services. Surprisingly, in the last decade, the nursing shortage has been acknowledged as a global issue, and predictions foresee a reduction of the nursing workforce (Oulton, 2006; OECD, 2013). This trend has been explained by a variety of factors, including the aging of the nurses, unfriendly work environments, and the reduced tendency of young generations to become nursing professionals (Oulton, 2006; De Oliveira et al., 2017).

Under this scenario, the intention to leave nursing has also been studied. Occupational and organizational factors, as well as physical burdens, long working hours, patient number assignment, psychological strain, and the effect that some of these factors have on the personnel health, have shown to contribute to this phenomenon (De Oliveira et al., 2017; Fügener et al., 2018). Particularly in Brazil, these labor-related factors have an important role in the intention to leave the profession, in a context where nurses besides health care assume management, administrative, and planning roles (De Oliveira et al., 2017). In this country, in the year 2017, there were 1.5 nurses per 1000 people while the mean established by the Organization for Economic Co-operation and Development Countries was 8.8 (OECD, 2019).

Facing the current and future nurse shortage exhibits the need to encourage the nursing profession's appeal, an aim that can be fostered by reducing the workers' physical, mental, and time pressures, thus increasing their job satisfaction and decreasing early retirements. Changing the way shifts are scheduled can contribute to this purpose while simultaneously minimizing costs for health care institutions(Fügener et al., 2018).

Along with nurse shortage, health care administrators face the need to deal with demand variations by coordinating available resources. To do so, demand upgrades and staffing flexibility are used to coordinate internal (nursing staff and beds) and external resources (contract nurses) (Gnanlet and Gilland, 2009). How these resources are to be coordinated will directly affect the quality of care and the overall costs, making the study of flexibility in healthcare a relevant topic. In the following sections, we will first review some important works in the nurse scheduling problem (NSP), study flexibility in general, and later focus on the study of flexibility in the health care sector, particularly in the way it has been addressed in the nurse scheduling problem.

2.2 Basic concepts

Although sometimes they go unnoticed, schedules are part of our daily life. We plan from simple to complex activities or events in given places and times. They also have great importance in business and industry, where producing, providing a service, or just performing a task or activity on time, may affect costs and customers' perception.

In the healthcare sector, scheduling (interchangeably used with rostering) is a critical task. It helps to manage existing resources like personnel, beds, and physical spaces which directly impact service quality and operational costs. For example, suppose fewer nurses than required at a given time of the day are rostered in a hospital unit. In that case, there will not be enough nurses to attend all patients, or they will be attended but have less time dedicated to them. In this example, creating a schedule for nurses that considers the unit's hourly demand of nurses can help to avoid being understaffed through the day. Furthermore, along with the demand of a unit, other aspects might have to be considered when scheduling nurses, like the minimum resting time between working hours, maximum working hours, nurses preferences, among others. This type of problem is known as the Nurse scheduling problem, and it is widely known for its practicality and complexity.

In Figure 2.1 a representation of a nurse schedule for a single unit is presented. In this case shifts S1 and S2 are allocated to four nurses along a 30 days planning horizon. Although not perceived in the figure, the allocation has to be made made under the consideration of a set of constraints like a maximum of one shift per day by nurse.

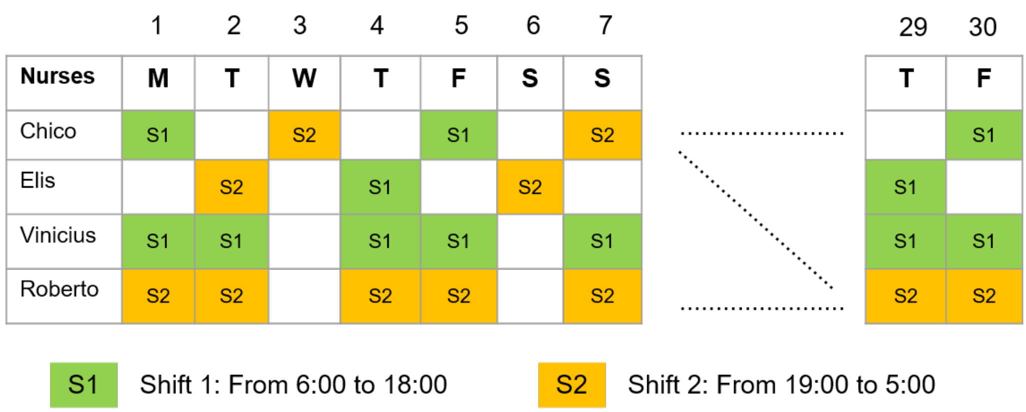

Figure 2.1: Illustration of a nurse schedule. Note: elaborated by the authors

The complexity of the NRP increases as healthcare institutions deal with supply and demand variations, this is, changes in the availability and need of healthcare services. One way in which managers deal with those changes is by introducing staff flexibility, through the use of agency, float pool and cross-trained nurses. Fügener et al. (2018) described each of these groups as follows:

- *Agency nurses*: External nurses who may be temporary employed at short notice to cover demand peaks. Usually associated with higher costs.

- *Float pool:* Nurses who are not dedicated to a specific unit (e.g hospitalization, cardiovascular care, surgery, intensive care, among others) inside a healthcare institution, but trained to cover all units connected with the pool and scheduled to a fixed shift pattern. Their assignment to a unit depends on the urgency needs of a unit for an additional nurse at the beginning of the nurse's shift.

- *Cross-trained nurses:* Internal nurses with a dedicated unit trained to cover one or more additional unit in case of demand peaks. The dedicated unit is denoted as "home unit" and the additional units as "float units".

In this work, we will focus on flexibility through the use of cross-trained nurses. As previously mentioned, cross-trained nurses are trained to work in units different from their dedicated ones. This selection is justified, as works considering cross-training in healthcare are very limited in literature, and there are only a few who analyze cross-training (CT) in mid-term nurse scheduling problems. All of those works will be further explored in Chapter 3. Furthermore, in practice, cross-training can provide better quality of care than temporary (from external agencies) nurses, reduce turnover, have a positive impact in morale, and reduce costs Inman et al. (2005).

To better understand the concept of cross-training, lets take a look at Figure 2.2. In the left side of the figure we see home units (large rectangles), each one of named after a number, in this case units 1, 2 and 3. As previously mentioned, these units might correspond to different areas or specializations of a healthcare institution, like surgery, pediatric care, among others. Note that for cross-training to be feasible, these units have to be different, but also sufficiently similar (Fügener et al., 2018). Each unit has dedicated nurses, also represented by numbers (in small rectangles), i.e unit 1 has ten dedicated nurses (N1 to N10), unit 2 has five (N11 to N15) and unit 3 also has five nurses (N16 to N20). Next to the home units are the float units, and initially no links exist between home and float units as no nurse is cross-trained to work in a unit different from its dedicated one. When wanting to represent that nurses are cross-trained from one unit to another we use the notation showed in the right part of the Figure 2.2, where the arrow links represent that certain nurses, or a percentage of nurses in unit 1 are cross-trained to unit 2. Dashed lines represent nurses that can work in float units with certain efficiency, usually lower than 100%.

Note that training links in configurations along this work will be represented between different units, and we will omit the link that exists between a unit and itself, as by default a unit can use its dedicated nurses. It is important to remember this to better understand the concept of a chain, that will be introduced later on.

To define the extend of cross-training three indicators are used, they can be simply defined as follows (Fügener et al., 2018):

- *CT Intensity*: The number (or percentage) of cross-trained nurses within a unit.
- *CT breadth*: The number of float units a cross-trained nurse is applicable to. The maximum breadth is determined by the number of units, if there are $|J|$ units in total, the maximum breadth is $|J| - 1$.
- *CT depth* : The level of productivity and quality of care, which cross-trained nurses are able to provide when working in a float unit.

Some examples of configurations with different CT indicators are depicted in Figure 2.3. In (a) we see a configuration with Intensity = 1, Breadth=1 and Depth=100%. This means one nurse from

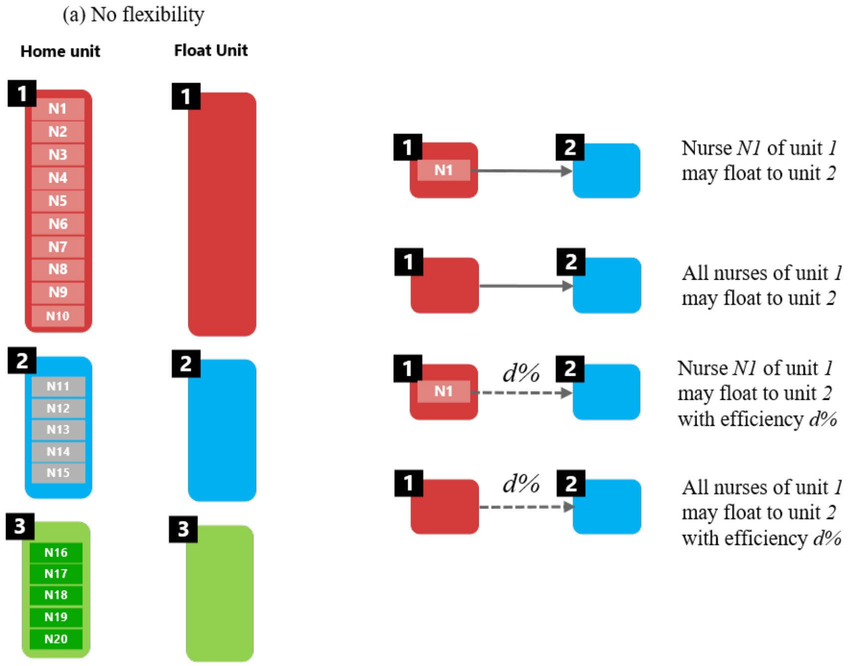

Figure 2.2: Generic illustration of cross-training policies.

each unit (intensity) is cross-trained to only one another unit (breadth) and can work in it with a 100% efficiency (depth). A variation of (a) is presented in (b) to illustrate a change on intensity, and in (c) to represent a change in breadth, as cross-trained nurses are allowed to work in another two different units from their home unit. Finally, (d) represents a case where the efficiency of the cross-trained nurses is 50% when they have to work in their float units, as identified by the dashed links.

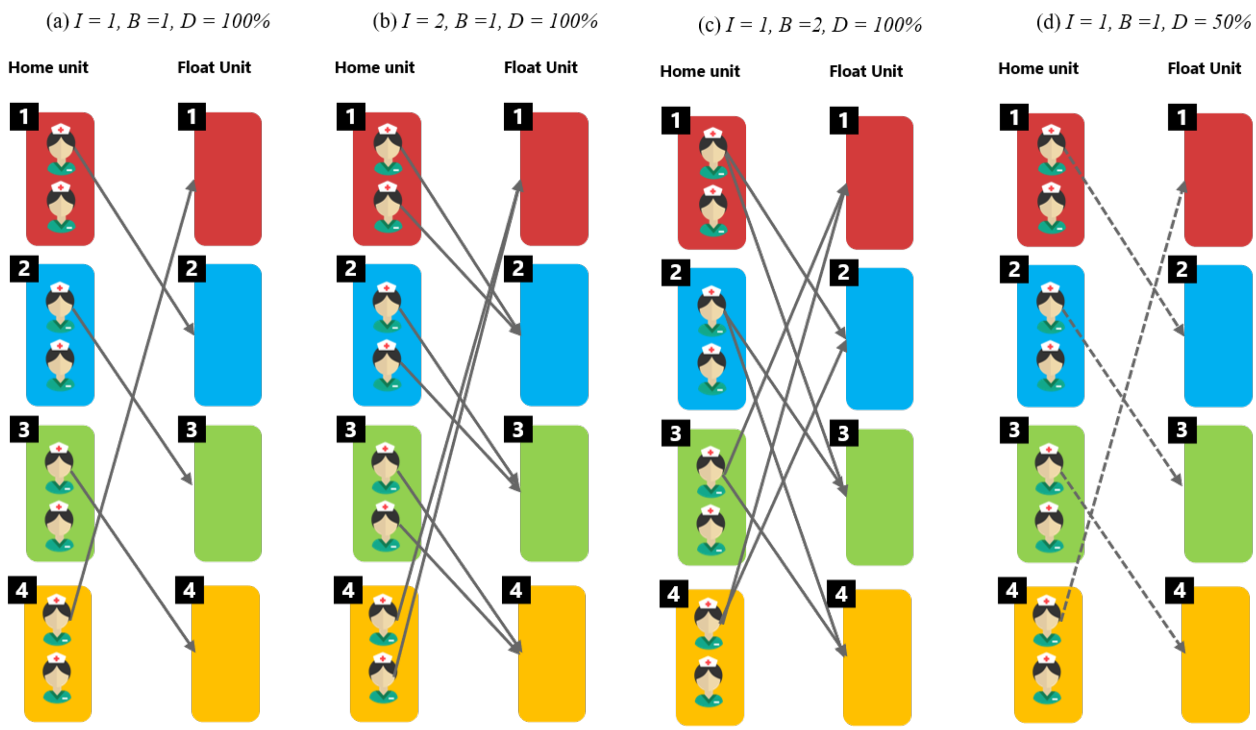

Figure 2.3: Examples of different indicators in cross-training configurations

The reader might note the assignments to float units in Figure 2.3 are arranged in a certain manner. This is because in this example the cross-trained nurses were assigned to float units according a specific *chaining* policy. We could define a CT policy (or configuration) as a principle or a way to arrange cross-trained workers to float units. Some of the most famous policies described by Inman et al. (2004, 2005), and the *one-for-each* policy proposed by Fügener et al. (2018) are represented in Figure 2.4. A brief description of these policies as provided by the same authors is presented below.

- *No cross-training*: Every nurse is dedicated to single unit and not allowed to work in any other. This configuration is depicted in Figure 2.2.

- *Reciprocal Pairs*: A CT policy where units are matched in pairs of two as illustrated in Figure 2.4 (a). Cross-training is allowed within the pairs only and an even number of units is required. This policy can also be found as a clustered policy.

- *Chaining*: A CT policy where trainings are arranged forming a chain. A *chain* can be described as a connected graph (Jordan and Graves, 1995a), and within a chain a path can be traced from home units to float units through assignment links. For example, in Figure 2.4 (b) we can trace a path from Home unit 1 to float unit 4, though the links 1-2, 2-3 and 3-4.

  According to Jordan et al. (2004), if we analyze this configuration as a graph, a chain is said to be *complete* if it is possible to start at any node and traverse the entire graph and return to the starting node without using any arc more than one time. A complete chain is *minimal* if it uses the minimum number of links (trainings in our case) such that the chain is complete. Given n units, there are $(n-1)!$ minimal complete chains.

- *n-to-all*: A CT policy where each unit provides n nurses that are cross-trained for all other units. As an example in Figure 2.4 (c) a single nurse form each unit is trained to work in the remaining units.

- *One-for-each*: A CT policy where $|J|-1$ nurses from each unit are trained, so there is one to work in each other unit. This policy is observed in Figure 2.4 (d).

- *Total cross-training*: This is a full flexibility scenario where all nurses in every unit are cross-trained to work in every other unit. This is equivalent to have a big pool of nurses. This policy is observed in Figure 2.4 (e).

The number of trainings for each policy changes according to the intensity and breadth indicators. For some policies one or both indicators are fixed, for example policies like reciprocal pairs and one-for-each have fixed breadth of 1, and the n-to-all policy has breadth $|J|-1$. This means that depending on the policy certain combinations of intensity and breadth are not possible. To exemplify this and the calculation of the number of trainings by policy, lets suppose we have 4 units, each with 5 nurses and want a policy with Intensity = 2 and Breadth = 2. Table 2.1 shows the number of trainings for each policy for this example, in bold the indicators that had to be fixed by the structure of the policy. This causes the number of cross-trainings by policy to be different when talking about the same combination of intensity and breadth. A full illustration of this case is shown in Figure 2.5.

In this work, we analyze the benefits of CT flexibility of each of the previous policies in the mid-term nurse scheduling problem. We do so, parting from a base model from the literature and

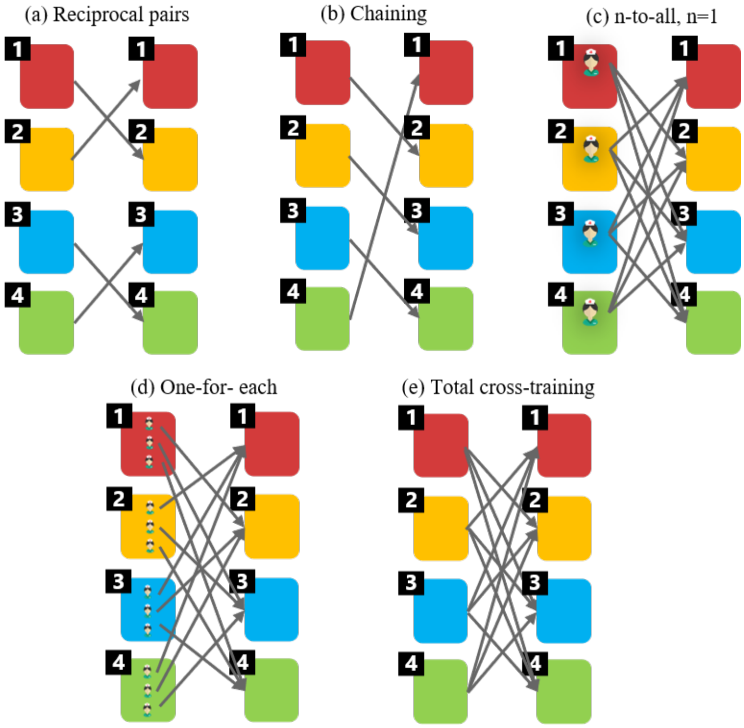

Figure 2.4: Illustration of common cross-training policies for a unit example

Table 2.1: Cross-trainings by policy in an illustrative example

Policy	Intensity	Breadth	# Trainings
Reciprocal pairs	2	**1**	8
Chaining	2	2	16
n-to-all	2	**3**	24
One-for-each	**3**	**1**	12
Total CT	5	3	60

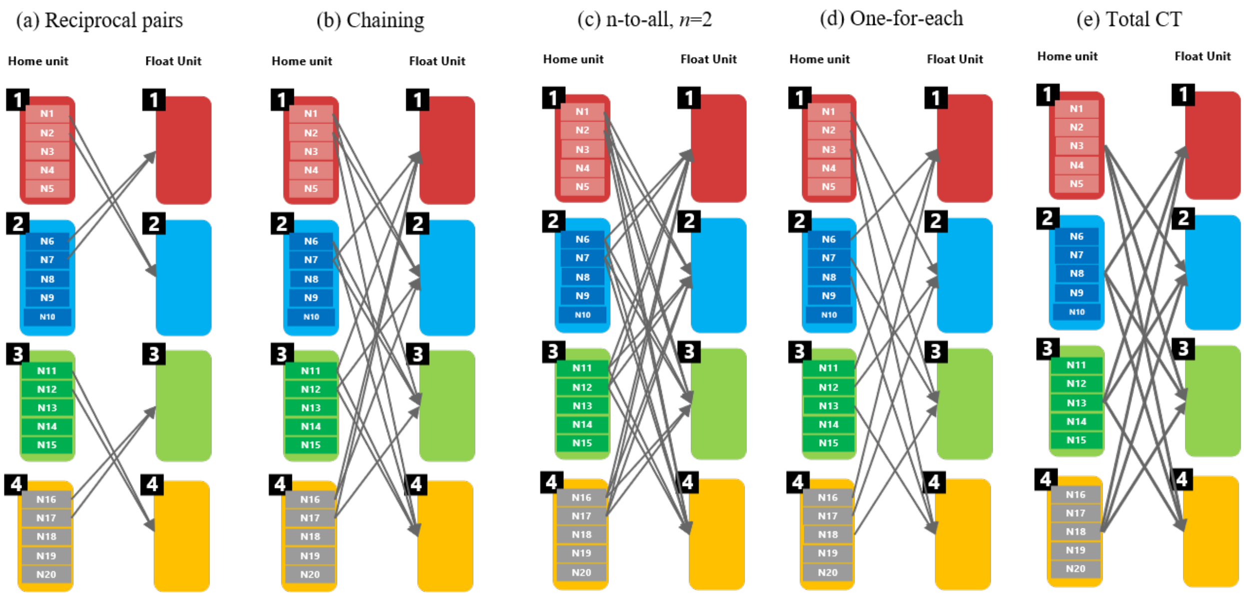

Figure 2.5: Illustration of policies in Table 2.1

proposing two new extensions. The first one considers the possibility to invest in flexibility and the second allows us to find the best chain configuration. A heuristic with the purpose to find the best chain is also proposed. Finally, a series of experiments are executed to help us to derive insights about CT flexibility, its indicators and the proposed models.

CHAPTER 3

Literature review

The scheduling or rostering of personnel is a common problem that health care institutions frequently face all over the world. This process consists of designing timetables for the staff to satisfy a given demand considering the workforce's skills, time, legal constraints, and employee preferences, among other factors (Ernst et al., 2004; Van den Bergh et al., 2013). The proportion that labor costs represent for many companies and the effect scheduling has over aspects like quality service and employees' job satisfaction have become a great motivation for studying this problem (Burke et al., 2004; Ernst et al., 2004). Furthermore, the known NP-Hard complexity of the problem (Solos et al., 2013; Hartog, 2016; Messelis et al., 2013; Rahimian et al., 2017) has led researchers to study throughout the years a wide variety of methods to solve different variations of the problem. For complete reviews on nurse scheduling the reader might refer to Ernst et al. (2004); Burke et al. (2004) and Van den Bergh et al. (2013), and for a full categorization of nurse scheduling problems De Causmaecker and Berghe (2011) provides a detailed description of their proposed notation.

We focus the literature review presented in this chapter on the works that approach flexibility in different contexts. First, we provide some background on the study of flexibility by discussing works that include flexibility at different decision-making levels in sectors distinct from healthcare. We particularly discuss flexibility in those sectors by focusing on types of flexibility, policies, multifunctionality, and other characteristics. Then, we discuss flexibility in healthcare in a similar manner while also remarking the similarities and differences with the previous sectors, and showing the main insights and findings that serve as a base for our research. Finally, our work's main contribution is stated.

3.1 A background on the study of flexibility

The industrial sector shifted from low-cost mass production of standard products to high-quality products and services with a certain personalization level in small batches that attend exclusive customer niches. With this change came the need to have more flexibility in manufacturing systems, this is, to have a higher ability to change rapidly, react or adjust to complexities, uncertainties, and transformations (Chauhan and Singh, 2011). In manufacturing, machines and workforce constitute the main resources, and when manufacturing flexibility can only be achieved with these two flexible resources, we talk about resource flexibility (Chauhan and Singh, 2011, 2014). Similarly, in the services sector, it is also possible to find resource flexibility represented by the workforce.

Qin et al. (2015) presented a wide review on *workforce flexibility* in operations management. They classified the workforce flexibility in literature by distinguishing five methods widely researched and implemented in practice:

- *Flexible working time*: creates flexibility in workforce capacity by permitting various labor hours, subject to constraints by-laws or consensus between employees and workers. This includes approaches like overtime, flexible workdays, annualized working hours, and working time accounts.

- *Floaters:* same as utility or floating workers are generalists that can be allocated to jobs when needed. They have the ability to perform a wider range of operations when compared to the regular workforce. Usually, floaters have a higher pay as a result of their generalist nature.

- *Cross-training:* In this method, workers obtain and maintain a series of skills that allow them to perform different tasks, so these workers can be assigned to a different unit when required. It requires choosing specific configurations like no cross-training, pooling, chaining, full cross-training, among others, establishing who is going to be trained to what task or unit. There are different costs and side effects of this practice, including reduced efficiency, productivity/quality loss, transfer costs, training costs, and additional wages.

- *Teamwork:* Functional or additive collaboration. This approach can have different advantages, like improving quality, reducing productivity loss, improving motivation, and reducing fatigue. Teamwork facilitates flexibility by the dynamic interaction of people with complementary and multiple skills. These multi-functional teams create flexibility by combining the necessary skills for high performance.

- *Temporary labor:* The contract of additional short-term labor. As for the name, they do not have a permanent contract with the organization and are used to cope with uncertain demands of products or services.

All these methods provide flexibility in workforce capacities, capabilities, or both. In addition to all of the above, Qin et al. (2015) remarked opportunities for research framed under three aspects: (a) under-researched issues of individual flexibility methods, where the author discussed that teamwork, as a method, has received less attention in the literature, due to the complexity to understand human behavior in dynamic team settings, and the need to extend the current research in the cross-training area; (b) problems from hybridizing of multiple flexibility methods, like the use of floaters, cross-training and teamwork to deal with changes in product mix; (c) the challenging use of workforce flexibility in complex systems where uncertainties are less predictable and changes are fast. In this work, we focus on the first aspect that calls for need of works extending the cross-training area.

Next, we further explore some relevant literature about the study of resource flexibility. We use the structure proposed in the works of Henao et al. (2015, 2019) who studied multiskilling in the retail industry. In Table 3.1 we organize the most relevant papers according to the following characteristics:

- *Resource Type (RT)*: Indicates if the paper focuses on machine flexibility(M), workforce (W) flexibility, or both.

- *Flexibility method (FM)*: When workforce flexibility is addressed, this characteristic describes which types of the workforce methods previously described are used: Flexible working time (FWT), floaters (F), cross-training (CT), teamwork (TW) or temporary labor (TL).
- *Decision Level of Humar resources (DL-HR)*: Abernathy et al. (1973) provided a hierarchy for workforce planning that has been cited in different studies (Campbell, 2011; Easton, 2011; Henao et al., 2015, 2019). These three levels of hierarchy, from higher to low, correspond to planning, scheduling and allocation, and can be helpful to analyze the human resource management process. Then, the DL-HR column in the table helps to identify which kind of decisions were addressed in each work. In the planning level the decision is to determine the staffing level (S) of each unit, department or area, usually on an annual or semiannual basis. Next, at the scheduling level the most common decisions are: (i) shift scheduling (SS) which consists on assigning workers to daily shifts, days off scheduling (DOS) assigning the days off between the working days, and tour scheduling (TS) that consists of doing the previous two decisions simultaneously. The planning horizon in scheduling goes from 1 to 4 weeks. Finally, the allocation phase (also referred as assignment sometimes), focuses on the day-to-day, changes and adjustments that need to be made between shifts across-departments, like the assignment of workers to tasks or units independently of their shifts.
- *Multifunctionality (MF)*: Indicates if in the study the employed model uses a parameter (Par) to fix the flexibility or if flexibility is a decision (Var).
- *Studied policies*: If applicable indicates the CT policies studied in the article. When configurations do not fit in any of the most widely known policies, like reciprocal pairs, chaining, full flexibility, all to all, n-to all, we classify these policies as custom policies. Also, when flexibility is a decision i.e. the model must decide what workers are trained to work in certain skills or units we say CT is a decision.
- *System type (ST):* if the proposed model is: (a) Homogeneous (Hom), where all task types, units, departments or resources in general are identical, meaning they have the same supply level, installed capacity (employees, machines etc.), identical demand, equal productivity, multiskilling or probability of absenteeism; heterogeneous (Het) meaning that not all task types, departments or factories are identical.
- *Uncertainty*: Indicates if the paper considered uncertainty in:
  - Demand(D): Indicates if variability of demand is allowed.
  - Capacity(C): Indicates if the problem considers the variability of capacity, this is if the resources can have uncertain speeds or productivity levels.
  - Supply(S): Indicates if the problem consider variability of supply or absenteeism.
- *Solution Method (SM)*: The solution methods used in the articles among them: Heuristics (H), Integer programming (IP), mixed integer programming (MIP), simulation (Sim), queuing theory (QT), Markov Process (MP), analytic (A), Robust Optimization (RO) and stochastic programming (SP).
- *Application (APP)*: The sector where the model is applied in a study e.g manufacturing (M), services (S), call centers (CC), health (H) and retail (R).

- *Planning horizon*: If scheduling is in the scope of the study this column indicates the schedule planning horizon.

- *Productivity modelling (PM)*: Shows the approach used in the model to represent the loss/gain productivity when individuals become multi-skilled or float to other units. As noted by Henao et al. (2015) different approaches can be found in the work of Yang (2007): (a) learning and forgetting (L/F), where individuals naturally learn, forget, and relearn, to eventually become more efficient in some skills; (b) productivity matrix (M) where workers have a productivity in each skill or unit, which is represented in a matrix; equal productivity (E), where workers are assumed equally productive in additional units/skills. Note that it is common for papers to use the terms efficiency and productivity interchangeably, but we will use the second term in this work.

It is worth nothing that some of the reviewed works consider flexibility with skills, meaning a dedicated worker has a specialized skill and can be trained to perform secondary skills as in Brusco and Johns (1998); Brusco (2008), in other cases flexibility is studied in a multi-department setting where workers in a unit, section or department have multiple skills and knowledge necessary to work in that unit, therefore cross-training them would mean to train them in the additional necessary skills and protocols of the additional destiny units. Articles with both approaches will be included in our review, but our work is applied to the second case.

In this section, we will focus on the studies with applications different from healthcare shown in Table 3.1. This table shows that most of the reviewed articles focus on workforce flexibility and that cross-training is the most common method to address it. Only the work of Mac-Vicar et al. (2017) uses additional workforce flexibility methods like flexible working time and temporal labor. A few of the reviewed works deal with machine flexibility, but as a common factor, all the reviewed articles refer to the principles on Jordan and Graves (1995b), which focused on machine flexibility, as their principles showed to be helpful when addressing flexibility in workforce. When looking at the decision level, the majority of the works cover the lowest level. In some cases, two decision levels were combined, and only the work of Easton (2011) integrates the whole decision hierarchy by staffing, scheduling, and allocating workers. Most of the authors who considered scheduling decisions chose to use short-term horizons of 1-day or 1-week.

Jordan and Graves (1995a) developed and provided insights on the benefits of flexibility, showing that it is possible to obtain most of the benefits of full flexibility with limited links, and this limited flexibility has the greatest benefits when products and plants are configured using a chain. In the work of Jordan et al. (2004) this chaining strategy was applied in cross-training, showing that chaining also yields most of the benefits of cross-training all workers. Furthermore, in their study, the complete chain previously defined, even when not optimal, showed to be robust, having a performance with low sensitivity to system utilization, how the chain is constructed, errors in parameters estimation, lack of data, and varying conditions. For the reasons mentioned above, when analyzing cross-training practices, a good proportion of the authors decided to use this policy in their work. Overall, this policy showed to provide good results when tested by different authors (Brusco and Johns, 1998; Brusco, 2008; Simchi-Levi and Wei, 2012; Hopp et al., 2004; Inman et al., 2004; Iravani et al., 2005a; Fiorotto et al., 2018), showing its value from a practical point of view. Finally, it is essential to mention that most of the studies, except the work of Fiorotto et al. (2018), ignore the order of the elements in the chain (products, workers). Strategically ordering the

Table 3.1: Reviewed articles on resource flexibility

Author, year	**RT**	**FM**	**DL-HR**	**MF**	**Studied policy**	**ST**	**Uncertainty (D/C/S)**	**Method**	**APP**	**Planning horizon**	**Productivity Mod**
Pinker and Shumsky (2000)	W	CT	S	Var	-	Hom, Het	Yes/Yes/No	QT + A	CC	-	L/F
Jordan and Graves (1995b)	M	-	A	Par	Chaining	Hom/Het	Yes/-/-	S + A	M	-	E
Brusco and Johns (1998)	W	CT	S + SS	Par	Chaining, Reciprocal pairs, Custom	Het	No/Yes/No	IP	S	1-day	M
Brusco (2008)	W	CT	A	Par	Chaining	Het	Yes/-/Yes	H + IP	S	-	M
Vairaktarakis and Winch (1999)	W	CT	A	Par	Custom	Het	Yes/-/No	IP + H	M	1-day	E
Agnihothri and Mishra (2004)	W	CT	S	Var	Chaining	Het	Yes/No/No	QT + S	S	-	M
Hopp et al. (2004)	W	CT	A	Var	Chaining	Het	Yes/Yes/No	IP + S	M	-	E
Inman et al. (2004)	W	CT	S + A	Par	Chaining	Hom	No/No/Yes	S	M	-	E
Jordan et al. (2004)	W	CT	A	Par	Full CT, Chaining, custom	Hom/Het	Yes/Yes/No	QT + S	M	-	E
Iravani et al. (2005b)	M, W	CT	A	Par	Chaining, Reciprocal pairs, custom	Het	Yes/Yes/No	S	CC	-	E
Campbell (2011)	W	CT	TS + A	Par	Custom	Het	Yes/No/No	H	S	1 month, 1-day	M
Easton (2011)	W	CT	S + TS + A	Var	Full CT, CT. As Decision	Het	Yes/Yes/Yes	MIP + S	S	-	M
Simchi-Levi and Wei (2012)	M	-	A	Par	Chaining	Hom	Yes/No/-	A	M	-	E
Andradóttir et al. (2013)	W	CT	A	Par	Chaining, Custom	Hom/Het	Yes/Yes/No	QT +A	S	-	E
Büke et al. (2016)	W	CT	A	Var	CT. as a Decision	Hom/Het	Yes/Yes/No	A	S	-	M
Olivella and Nembhard (2016)	W	CT	A	Var	CT. as a Decision	Hom	Yes/No/Yes	H	M	-	E
Henao et al. (2015)	W	CT	TS + A	Var	CT. as a Decision	Het	No/No/No	MIP	R	1-week	E
Henao et al. (2016)	W	CT	A	Var	Chaining	Het	Yes/No/No	MIP+ RO	R	1-week	E
Mac-Vicar et al. (2017)	W	FWT, TL, CT	SS + A	Par	Custom	Hom	Yes/No/No	H	R	1-day	E
Henao et al. (2019)	W	CT	S	Var	Chaining	Hom	Yes/No/No	A + H + S +IP	R	1-week	E
Taskiran and Zhang (2017)	W	CT	TS + A	Var	CT. as a Decision	Het	No/Yes/No	IP	CC	1-week	M
Fiorotto et al. (2018)	M	-	A	Var	Chaining, Clustered	Het	Yes/Yes/No	IP	M	-	E

chain can be more effective, particularly when systems are heterogeneous, which can cause some permutations to perform better than others (Inman et al., 2005; Paul and MacDonald, 2014).

Almost an equal number of papers considered flexibility as a decision variable and as a parameter. The most common approaches of the papers where flexibility was a decision were: (a) to enforce a policy e.g chaining or any of its variations and, in the case of CT flexibility, deciding how many employees in each skill or unit to train, like in the work of Henao et al. (2016); (b) let the model decide the amount of flexibility and also in which way is applied (Henao et al., 2015; Taskiran and Zhang, 2017; Olivella and Nembhard, 2016; Easton, 2011; Büke et al., 2016). Also, most of the articles considered heterogeneous systems, where the resources can have different capacities, productivity, among other properties. When looking at uncertainty, most articles dealt with uncertainty in demand, capacity, or both; fewer articles considered uncertainty in supply, and only Easton (2011) managed to include the three types of uncertainty.

Most of the reviewed studies that work with cross-training propose models where productivity, also referenced in the literature as cross-training *depth*, is equal among cross-trained and dedicated workers. However, as Pinker and Shumsky (2000) and other authors perceived, workers are not as productive in cross-trained tasks (units, or departments) as they are in their original or specialized tasks, which translates in quality loss. To capture this effect, the most used method was to use a productivity matrix where productivity in additional skills was reduced i.e a worker had 100% efficiency while performing its original task or while working in its primary department or unit, and a reduced productivity when working in secondary tasks or units. Regarding these effects Taskiran and Zhang (2017) found out that as the efficiency in secondary skills decreases, the benefits of cross-training could quickly fade away, and that adding more skills could result in an additional loss of efficiency. In their work Brusco and Johns (1998) concluded that a 50% productivity was sufficient to attain most of the cost savings available for cross-training at a 100% productivity. Finally, in the work of Easton (2011), when the productivity of cross-trained workers was less than 100%, a greater cross-training forced a trade-off between workforce size and capacity shortages in a services environment.

Regarding solution methods, ours review shows that MIP and IP models are preferred by authors, sometimes in combination with techniques like simulation, heuristics and analytic methods. Queuing theory and simulation where also on a wide range of the works, being useful when workers were modeled as servers that perform a set of tasks in service environments. Manufacturing and services were the groups with more participation in the reviewed articles, followed by the retail sector where there are only a few works in the last consecutive years, all of them by Henao et al. (2019); MacVicar et al. (2017); Henao et al. (2016, 2015) who studied different chaining configurations, and proposed models and methods to manage multiskilled workforce in the context of a retail industry.

3.2 Flexibility in the healthcare sector

Just as it happens in manufacturing, it is possible to identify resource flexibilities in the healthcare sector, and managing them can help to improve response to uncertainties. According to Gnanlet and Gilland (2009), there are two main types of resource flexibilities that managers deal with within healthcare, namely demand upgrades and staffing flexibility, that are usually used to coordinate staff and beds (internal resources) and contract nurses (external resources). As explained by the same authors, with demand upgrades, if there is not availability of beds for patients in less critical units, they are promoted to a higher critical unit if possible. With staffing flexibility, cross-trained nurses, i.e., educated to labor in other work units apart from their dedicated ones, float and contract (temporal labor) nurses are used. As a take from these authors, their work allowed them to conclude that these two types of flexibility are complementary, although staffing flexibility has on average more significant benefits. For this work we will focus on staffing flexibility using cross-training as a method. The present work focuses on staffing flexibility.

The most relevant articles dealing with flexibility in healthcare are organized in Table 3.2 according to the same characteristics as Table 3.1. The table reveals that most authors prefer to work with staffing flexibility (i.e workforce flexibility), as demand upgrades (DU) were rarely addressed. Also, in healthcare, we see that it is more common to work with a combination of flexibility methods, as more of the reviewed articles use at least two, and Wright and Bretthauer (2010) use four different flexibility strategies. Cross-training was the most used method as it showed to provide good benefits.

In over half of the articles, the authors combine at least two decision levels, being the lowest level (allocation) covered in almost all cases. Only the works of Gnanlet and Gilland (2009) and Maenhout and Vanhoucke (2013) included the whole hierarchy in the same work. Of the six works that consider flexibility in their scheduling models, short (1-day) and midterm schedules (e.g., monthly) were the most common planning horizons. As solution methods, most of the studies used IP models and stochastic programming. Unlike papers in the previous section, simulation was used only in one article, and fewer papers combined solution methods. Heterogeneous systems were considered in all but one study, as was uncertainty in demand, and uncertainty in supply related to absenteeism was considered in only two articles. A little less than half of the works considered uncertainty in capacity.

Regarding policies, chaining was found to provide the most significant benefits in the work of Gnanlet and Gilland (2014) under centralized and decentralized decision-making. Paul and MacDonald (2014) highlighted the need to consider this policy to make resource allocation decisions, mainly when the budget is very limited. Likewise, Inman et al. (2005) remarked the reliability of this policy to provide benefits at a fraction of the full cross-training, its good efficiency with a realistic absenteeism rate, and better performance than the reciprocal pairs policy. Fügener et al. (2018) did not find a significant difference between the performance of chaining and reciprocal pairs but proposed a new one-for-each policy that showed to yield superior solutions than the most common policies, including chaining, n-to-all and reciprocal pairs.

Inman et al. (2005); Paul and MacDonald (2014) and Jordan et al. (2004) noted that there are different ways to structure policies like reciprocal pairs or chaining. How these cross-training policies are structured gains relevance, particularly in heterogeneous systems (with different numbers of nurses, absenteeism, patient census), where not all permutations will perform the same. Despite being acknowledged in some works, and most of the articles considering heterogeneous systems, this previous aspect was not analyzed in any of the reviewed works. Finally, four studies decided to consider flexibility through cross-training as a decision, two of which forced existing policies and the articles of Gnanlet and Gilland (2009) and Wright and Mahar (2013) did not force any policy, letting the model decide the way cross-training flexibility was applied.

In the related literature regarding flexibility in healthcare, the authors derived some conclusions about the cross-training indicators introduced in Section 2.2. Regarding breadth and intensity, in Wright and Bretthauer (2010) the best improvement occurred when nurses were trained to work in

Table 3.2: Reviewed articles on resource flexibility in healthcare

Author, year	RT	FM	DL-HR	MF	Studied policy	ST	Uncertainty (D/C/S)	Method	APP	Schedule Planning horizon	Productivity Mod
Trivedi and Warner (1976)	W	F	A	Par	-	Het	Yes/No/No	IP	H	-	E
Gnanlet and Gilland (2009)	DU, W	CT, TL	S + SS + A	Var	CT. As Decision	Het	Yes/No/No	SP	H	1 day, 1 month	E
Inman et al. (2005)	W	CT	A	Par	NCT, Reciprocal Pairs, Chaining, All fro all, Total CT	Het	Yes/No/Yes	S + A	H	-	E
Maenhout and Vanhoucke (2013)	W	CT	S+SS+A	Par	-	Het	No/No/No	IP+DP	H	4-week	E
Gnanlet and Gilland (2014)	W	CT, TL	S + A	Var	Full CT, partial, chaining	Het	Yes/Yes/No	SP	H	-	M
Paul and MacDonald (2014)	W	CT, TL	S + A	Var	Chaining	Het	Yes/Yes/No	MIP +H +A	H	-	E,M
Li and King (1999)	W	CT, FWT	S	Var	-	Het	Yes/Yes/No	GP	H	-	M
Bard and Purnomo (2005)	W	FWT, F, TL	SS + A	Par	-	Het	No/No/Yes	IP	H	1-day	E
Wright and Bretthauer (2010)	W	FWT, TL, F, CT	SS + A	Par	-	Het	Yes/Yes/No	IP	H	5-weeks	M
Wright and Mahar (2013)	W	F	SS	Var	-	Hom	Yes/No/No	IP	H	5-weeks	E
Fügener et al. (2018)	W	CT	SS	Par	Chaining, one-for-all, n-to-all, reciprocal pairs, No flexibility	Het	Yes/No/No	IP	H	14-days	E

one additional unit. With the addition of extra units, the benefits tended to be less representative. These results were confirmed with their findings in one of their later works, where despite increasing benefits when going from 1 to 3 additional units to work, the returns diminished while cross-utilization increased (Wright and Mahar, 2013). Fügener et al. (2018) reports a similar conclusion, stating that increasing cross-training intensity and breadth leads to more savings, but the growth of benefits decreases with the increase of cross-training, and intensity lead to stronger results when compared to breadth. In another study Gnanlet and Gilland (2014) highlights that as cross-traininig breadth increases, cross-trained workers might be less cost-efficient than specialists.

As mentioned in different works from the healthcare sector and in others, cross-training can affect service quality. As noted by Maenhout and Vanhoucke (2013) a delicate trade-off exists as the higher the flexibility the higher are the cost efficiency and job satisfaction, but the lower the effectiveness of providing high-quality care. In this case, cross-trained nurses could not be as productive as regular or dedicated nurses, given their inexperience with protocols in new tasks or units (Paul and MacDonald, 2014) which will affect the quality of care. Therefore, a few authors discuss the effect of cross-training depth, like Gnanlet and Gilland (2014) who worked with nurse staffing, and noted that there is a productivity threshold after which the number of cross-trained nurses needed to benefit from flexibility reduces, and when the cost of training increases this threshold increases. In contrast, when CT cost is low, this threshold is low, and the number of CT nurses increases. In another work dealing with staffing Paul and MacDonald (2014) analyze the impact of quality of cross-trained staff and found that cross-training was beneficial when the quality of CT was above 0.7 (70% depth). These results contrast with other papers in the service industry where a 50% depth was enough to reach most of the savings of a completely flexible scenario (Brusco and Johns, 1998; Brusco et al., 1998).

Main contributions

Despite the benefits that adding flexibility through cross-training can bring to the quality of schedules, as seen in Table 3.2 and in the previous discussions, in the last few years, only a few studies include this type of flexibility in their mid-term nurse scheduling models. From the reviewed articles, we observed that: (a) when using chaining as a policy, all the articles ignore the order of units and do not worry about finding the best chain configuration for the problem. Finding the best permutation, as stated by different authors, could be of relevance when systems are heterogeneous; (b) none of the articles proposes a model that contemplates the possibility to invest in flexibility while simultaneously finding the best way to apply it, and there was only one article that considered cross-training as a decision variable while scheduling. The previous findings motivate us to extend the work of Fügener et al. (2018), one of the most recent works that includes cross-training in the mid-term nurse scheduling problem. In their work, flexibility is considered a parameter, e.g. given a flexibility policy, it is known how many nurses will be trained in each area (Intensity) and to which unit (or units), leaving out the possibility of considering training as an investment. In this work, we extend their model to one where it is possible to invest in flexibility and analyze two ways to improve *chaining*, a policy that has shown to provide good results in previous works, by creating a heuristic to create the chain and with a model that produces the best chain. Then, the models look forward to find the cross-training configuration and schedule for every unit such that the understaffing and overtime costs are minimum, subject to a series demand and time related constraints.

CHAPTER 4

Mathematical formulations for the mid-term nurse scheduling problem

4.1 Base model

As shown in the different works from the literature review, the nurse scheduling problem has been widely studied over the years. Consequently, different variants and models have been proposed, with a variety of approaches, constraints, and parameters. To develop our analysis, we take as base the model proposed by Fügener et al. (2018) who were the first to develop a mid-term scheduling model including cross-training policies. We will extend their work by proposing and analyzing some variations of the model, one that allows flexibility as a decision and one that finds the best chain. The model sets, parameters, and variables of the base model proposed by Fügener et al. (2018) are shown in Tables 4.1 to 4.3, the objective is represented by expression (4.1), and constraints from (4.2) to (4.12).

We consider a set of nurses I, each of which has an assigned home unit HU_i where each nurse i is dedicated. Each nurse who belongs to a unit in J has to be scheduled throughout a horizon of $|D|$ days, formed by $|P|$ periods each. The periods in a day serve to describe the set of shifts S each with a length L_s and respective first and last working periods Beg_s and End_s. A day is divided into $|P|$ periods, and each shift is formed by a number of consecutive periods. So for example a shift 1 can cover from period 1 to period 12, and a shift 2 from period 13 to 24 for two 12 hour shifts. Note that the shifts can overlap. Nurses can not work more than one daily shift, and once assigned to a unit, they have to stay a number of consecutive $Stay^{min}$ days in it, with a maximum of Str^{max} consecutive working days. There is a minimum rest time R^{min} between shifts, a maximum number of working W_i and overtime hours OT_i^{max} that a nurse can work in the planning horizon. The flexibility configuration works as a parameter in the base model. This configuration is imposed by the use of the I_j parameter, which contains the nurses that are allowed (because they are either trained, or dedicated) to work in unit j.

For every period, day and unit there is a minimum demand Q_{jdp}^{min} and a a target demand Q_{jdp}^{NPR} of nurses. When those demands are not met it is said to be a number of missing nurses to cover the minimum and target demands, giving respective place to the Δ_{jdp}^{min} and Δ_{jdp}^{NPR} variables. In addition, variables y_{isjd} and ot_i respectively indicate if nurse i works on shift s, on unit j and day d and the overtime hours of nurse i in the whole planning horizon.

The objective function of the model is expressed in terms of cost for undercoverage and overtime cost, where $C^{\Delta min}$, $C^{\Delta NPR}$ are the cost for missing the minimum and target demand of nurses and C^{ot} is the overtime cost per hour. The authors assign weighted costs such that $C^{\Delta min} > C^{\Delta NPR} > C^{ot}$ but make clear that these parameters might differ according to specific needs and characteristics of hospitals.

Table 4.1: Indices and sets

Indices		**Sets**	
j	index for units	J	set of units
i	index for nurses	I	set of nurses
s	index for shifts	I_j	subset of nurses who may work in unit j
d	index for days	S	set of shifts
p	index for periods in a day	D	set of days
		P	set of periods

Table 4.2: Model Parameters

Parameter	**Description**
A_{sp}	1, if shift s covers period p;
Beg_s	First working period in shift s;
End_s	Last working period in shift s;
L_s	Length of shift s (in hours);
R_{min}	Minimum rest time between shifts (in hours);
Q_{jdp}^{min}	Minimum demand for nurses in unit j on day d in period p;
Q_{jdp}^{NPR}	Target demand in unit j on day d in period p with higher nurse-to-patient ratio;
HU_i	Home unit of nurse i;
H	Minimum proportion of working time in home unit;
$C^{\Delta min}$	Costs of each nurse missing to cover minimum demand ($/nurse);
$C^{\Delta NPR}$	Costs of each nurse missing to satisfy target demand ($/nurse);
C^{ot}	Costs for each overtime hour ($/hour);
$Stay^{min}$	Minimum number of consecutive days a nurse is assigned to a particular unit;
Str^{max}	Maximum stretch of consecutive on duty working days;
W_i	Total regular working hours of nurse i in the planning horizon;
OT_i^{max}	Maximum overtime hours of nurse i in the planning horizon;
y_{isj0}	1, if nurse i works shift s in unit j on the last day of the previous planning horizon, 0 otherwise

Table 4.3: Variables

Variable	Description
y_{isjd}	1, if nurse i works shift s in unit j on day d, 0 otherwise
Δ_{jdp}^{min}	Number of nurses missing to cover the minimum demand in unit j on day d in period p
Δ_{jdp}^{NPR}	Number of nurses in addition to minimum demand missing target demand Q_{jdp}^{NPR} in unit j on day d in period p
ot_i	Overtime hours of nurse i over the whole planning horizon

Model

$$Minimize \quad \sum_{d\in D}\sum_{p\in P}\sum_{j\in J}\left(C^{\Delta min}\Delta_{jdp}^{min} + C^{\Delta NPR}\Delta_{jdp}^{NPR}\right) + \sum_{i\in I} C^{ot} ot_i \tag{4.1}$$

Subject to:

$$\sum_{s\in S}\sum_{i\in I} A_{sp} y_{isjd} + \Delta_{jdp}^{min} \geq Q_{jdp}^{\min} \quad \forall j \in J, d \in D, p \in P \tag{4.2}$$

$$\sum_{s\in S}\sum_{i\in I} A_{sp} y_{isjd} + \Delta_{jdp}^{min} + \Delta_{jdp}^{NPR} \geq Q_{jdp}^{NPR} \quad \forall j \in J, d \in D, p \in P \tag{4.3}$$

$$y_{isjd} \leq 0 \qquad \forall \quad j \in J, i \in I \backslash I_j, s \in S, d \in D \tag{4.4}$$

$$\sum_{s\in S}\sum_{j\in J} y_{isjd} \leq 1 \qquad \forall \quad i \in I, d \in D \tag{4.5}$$

$$\sum_{s\in S}\sum_{j\in J}\sum_{d\in D} L_s y_{isjd} - ot_i \leq W_i, \quad \forall i \in I \tag{4.6}$$

$$ot_i \leq OT^{max} \quad \forall i \in I \tag{4.7}$$

$$\sum_{s\in S}\sum_{d\in D} L_s y_{is(HU_i)d} \geq H^* W_i \quad \forall i \in I \tag{4.8}$$

$$\begin{aligned} |P| - \sum_{s\in S}\sum_{j\in J} End_s y_{isjd} + \sum_{s\in S}\sum_{j\in J} (\mathrm{Beg}_s - 1)\, y_{isj(d+1)} \\ \geq R^{\min} \sum_{s\in S}\sum_{j\in J} y_{isj(d+1)} \quad \forall i \in I, d \in D \backslash \{|D|\} \end{aligned} \tag{4.9}$$

$$\sum_{s\in S}\sum_{j\in J} \sum_{d}^{\min(|D|;d+Str^{\max})} y_{isjd} \leq Str^m ax \quad \forall i \in I, d \in D \tag{4.10}$$

$$\begin{aligned} \sum_{s\in S} \sum_{\tau=d+1}^{\min\left(d+Stay^{min}-1;|D|\right)} y_{isj\tau} \geq \left(Stay^{min} - 1\right) * \sum_{s\in S} (y_{isjd} - y_{isjd-1}) \quad \forall i \in I, j \\ \in J, d \in D \cup \{0\} \end{aligned} \tag{4.11}$$

$$y_{isjd} \in \{0, 1\}, \Delta_{jdp}^{\min}, \Delta_{jdp}^{NPR}, ot_i \geq 0 \tag{4.12}$$

In their respective order, the objective and constraints in the previous model:

- (4.1) Correspond to the objective function that minimizes the weighted costs for undercoverage of minimum and target demand and of overtime.

- (4.2) Determine the undercoverage of minimum demand for each unit, day and period.
- (4.3) Measure the additional deviations from the target NPR adjusted demand.
- (4.4) Ensure that nurses are only assigned to units they are trained to.
- (4.5) Guarantee that each nurse may only work one shift in one unit per day.
- (4.6) Determines the overtime for the whole planning horizon for each nurse.
- (4.7) Limits the overtime hours for each nurse.
- (4.8) Ensure that each nurse works at least a pre-defined proportion o H of the regular working time in the home unit.
- (4.9) Establishes the minimum rest time between shifts on consecutive days.
- (4.10) Defines the maximum number of consecutive working days.
- (4.11) Force each nurse to stay in a particular unit and to work consecutively for at least $Stay^{min}$ days once scheduled to it.
- (4.12) Define the decision variables domains.

To obtain the daily demand the following expressions are considered, where E_{jdp} corresponds to the deterministic expected demand and u_{jd} is a number from a discrete uniform distribution $u_{jd} \sim U_{jd}\{-n; n\}$. These constraints consider deviation for expected demands during the night shift, and to simplify experiments, suppose that the target demand is only one nurse over the minimum demand.

$$Q_{jdp} = E_{jdp} + u_{jd} \qquad \forall j \in J, d \in D, \{p \in P \mid p \leq 16\} \tag{4.13}$$

$$Q_{jdp} = E_{jdp} \qquad \forall j \in J, d \in D, \{p \in P \mid p \geq 17\} \tag{4.14}$$

$$Q_{jdp}^{NPR} = Q_{jdp}^{min} + 1 \qquad \forall j \in J, d \in D, p \in P \tag{4.15}$$

Fügener et al. (2018) do not assume deviation for expected demands during night shifts (from period 17 onwards), which is why the expression (4.14) does not uses the u_{jd} to simulate the effect of stochastic demand and (4.13) accounts for this effect.

4.2 Model 1. Flexibility as a decision variable

For this model, we consider the possibility of investing in flexibility. Unlike the base model, where the flexibility configurations are fixed to analyze the benefits, this model creates the best distribution of cross-training links between units. So, flexibility is no longer fixed from the beginning, but there is a limiting budget or links (maximum number of trainings) and the purpose of the model is to find the best way to cross-train nurses while respecting this limit. With this objective in mind, we add constraints to the previous model that will consider the investment in flexibility or find the minimum number of links if required, such that the costs are minimized. A binary variable w_{ij} is introduced; this variable will take the value of 1 if nurse i is trained to work in unit j and 0 otherwise. In this model we re-define the set I_j as the subset of dedicated nurses in unit j, considering that each nurse can only have one dedicated unit. Then, by default $w_{ij} = 1$ if $i \in I_j$, this is, if nurse i is dedicated to unit j the nurse is allowed to work in j, and the model will decide to cross-train or not this nurse to work in other units.

We also include a parameter DC that represents the desired cross-trainings that the decision maker is willing to invest in.

The new mathematical formulation is given by objective in (4.1) and constraints (4.2)-(4.3), (4.5)-(4.12) used in the base model, and the following new constraints:

$$y_{isjd} - w_{ij} \leq 0 \quad \forall i \in I, s \in S, j \in J, d \in D \tag{4.16}$$

$$w_{ij} = 1 \quad \forall j \in J, i \in I_j \tag{4.17}$$

$$\sum_{i \in I} \sum_{j \in J} w_{ij} - |I| \leq DC \tag{4.18}$$

$$w_{ij} \in \{0, 1\} \tag{4.19}$$

The set of constraints in (4.16) does not allow for a nurse to be assigned a unit in a given day, unless it is trained to work in it. Constraints in (4.17) assign nurses to their respective home units, and constraints in (4.18) establish that the number of cross-trainings has to be less or equal to a number DC, which is the number of desired cross-trainings. This constraint can also help to force the use of a minimum or maximum number of cross-trainings, representing the investment in flexibility. If desired, by adding the objective of minimizing the number of links to the objective function, the minimum number of cross-trainings can also be found. The policy (h) in Figure 5.1, represents the solution for the example with CT as a decision, that minimizes both the costs in (4.1) and number of links.

4.3 Model 2. Finding the best chain

The principles and benefits of the chaining policy were previously studied by Jordan and Graves (1995b), showing that with chaining it was possible to obtain most of the benefits of full flexibility. In the work of Fügener et al. (2018) the chaining policy was analyzed for the mid-term nurse scheduling problem, showing good results. However, this policy was always forced in the same way, i.e. every unit cross-trains workers for the next unit, except for the last unit, which cross-trains nurses to the first one, with fixed positions for units. Doing an analysis similar to the one made by Fiorotto et al. (2018) for the lot-sizing problem, this might not correspond to the *best chain*. As mentioned before a chain can be seen as a connected graph, a minimal chain uses the minimum number of links such that the chain is complete and there are $(n-1)!$ minimal complete chains if there are n units. See, for example, in Figure 4.1 that by changing the units' positions on the left to right, we obtain a different chain that could have either a best or worse performance. Note that the changing of positions is a visual aid to help the reader to perceive the chain, but the positions might be left unchanged, and the links could join the same units forming a chain.

It is possible to find the best chain by adding constraints that allow every unit to cross-train nurses to additional units, however, this way of formulating does not stop the creation of sub-chains (closed clusters). To prevent this from happening, the following formulation is proposed. It is based on the idea that a minimal chain is in reality a long tour, where every unit has to be visited once from another unit. So, the tour $1-2-3-4-1$ is the chain represented in the left of Figure 4.1, where in addition to itself the unit provides nurses to the following unit, by exception of the last one which provides to the first unit. A different and valid chain could be formed by $4-1-3-2-4$ in the right of the same figure.

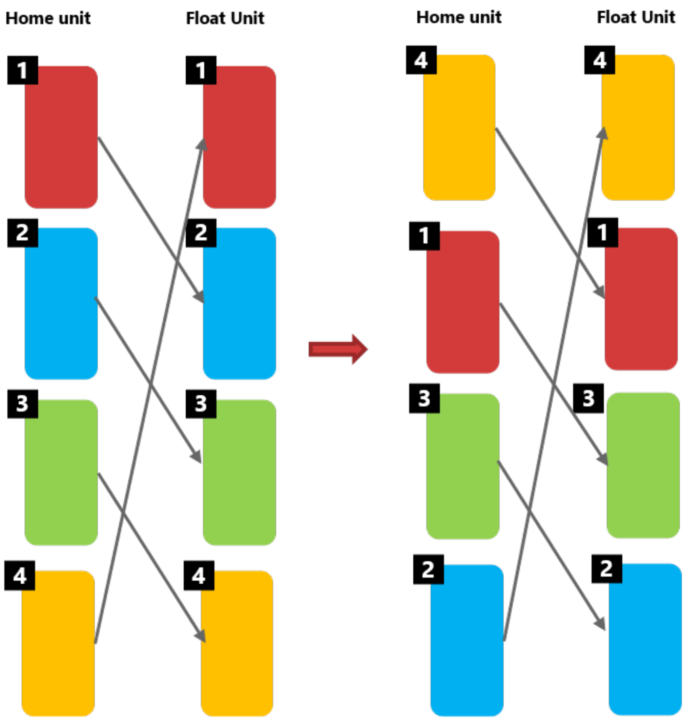

Figure 4.1: Different chaining configurations

To formulate this model the following additional subsets and parameters are included in the formulation:

- n_j is the number of nurses to be cross-trained within a unit j. When intensity is the same for all units and greater or equal to 1 we will represent this parameter only as n where $n = Intensity$. In contrast, when intensity is given as a proportion or percentage, e.g. if intensity $= 0.5$ (or 50%) meaning half of the nurses in each unit can float to other unit, this parameter is stated as $n_j = \lceil Intensity * |I_j| \rceil$.
- I_j^n subset of n nurses in set I_j who are to be cross-trained and may work in another unit or units. Note that in this case the parameter I_j must contain only the dedicated nurses of each unit.
- $\alpha_j = min\{n_l : \forall l \in J, l \neq j\}$ represents the minimum number of cross-trained nurses among the units different than j. For example, if units from 1 to 4 have 4, 2, 3, 5 respective number of nurses, $\alpha_1 = 2$.
- $\omega_j = max\{n_l : \forall l \in J, l \neq j\}$ represents the maximum number of cross-trained nurses among the units different than j. For example, if units from 1 to 4 have 4, 2, 3, 5 respective number of nurses, $\omega_1 = 5$.

Considering this, when $Intensity \geq 1$ in addition to objective (4.1) and constraints (4.2), (4.3), (4.5) - (4.12), (4.16), (4.17), and (4.19) we add the following constraints:

$$\sum_{i \in I} w_{ij} - |I_j| = n_j * Breadth, \quad \forall j \in J \tag{4.20}$$

$$\sum_{\substack{j \in J \\ j \neq j'}} w_{ij} = Breadth, \quad \forall i \in I^n_{j'}, \quad j' \in J \tag{4.21}$$

$$\sum_{\substack{j \in J \\ j \neq j'}} w_{ij} = 0, \quad \forall i \in I_{j'} \backslash I^n_{j'}, \quad j' \in J \tag{4.22}$$

$$w_{ij} - k_{(HU_i)j} \leq 0, \quad \forall i \in I, \quad j \in J \tag{4.23}$$

$$\sum_{i \in I_j} w_{ij'} \geq k_{jj'}, \quad \forall j' \in J, j \in J \tag{4.24}$$

$$\sum_{\substack{j \in J \\ j \neq j'}} z_{jj'} = 1, \quad \forall j' \in J \tag{4.25}$$

$$\sum_{\substack{j' \in J \\ j' \neq j}} z_{jj'} = 1, \quad \forall j \in J \tag{4.26}$$

$$z_{jj'} - k_{jj'} \leq 0, \quad \forall j \in J, j' \in J \tag{4.27}$$

$$u_1 = 1 \tag{4.28}$$

$$2 \leq u_j \leq |J| \quad \forall j \neq 1 \tag{4.29}$$

$$u_j - u_{j'} + 1 \leq (|J| - 1)\left(1 - z_{jj'}\right), \quad \forall j \neq 1, \forall j' \neq 1, j \neq j' \tag{4.30}$$

$$k_{jj'} \in \{0, 1\}, \quad \forall j' \in J, j \in J \tag{4.31}$$

$$z_{jj'} \in \{0, 1\}, \quad \forall j' \in J, j \in J \tag{4.32}$$

$$u_j \geq 0, \quad \forall j \in J \tag{4.33}$$

Constraints in (4.20) fix the number of nurses to be cross-trained in every unit, according to the desired intensity and breadth. In (4.21) and (4.22), only $n = Intensity$ nurses from each unit are cross-trained and may float to another units.

Constraints (4.23) create a link between the nurses trained to units and the k_{jj} variable. Basically, if a nurse i is trained to another unit j we say there is a link between the home unit of nurse i and the new float unit j. So for example, if $w_{12} = 1$ and $HU_1 = 1$ (the home unit of nurse 1 is unit 1) then $k_{12} = 1$. Constraints (4.24) express that if k_{jj} exists is because at least one nurse from unit j is cross-trained to work in unit j'. To force that across the assignments in the previous constraints a chain exists, we introduced a new variable $z_{jj'}$. This variable will guarantee that a chain exits with the help of constraints from (4.25) to (4.30).

Particularly, constraints (4.28) to (4.30) correspond to the well known Miller–Tucker–Zemlin (MTZ) subtour elimination constraints for the TSP, where a new variable u_j is added ($j = 1, ..., J$) to avoid the creation of sub-tours, and in this case sub-chains or clusters of chains. As a side note, a minimal chain will be found when breadth is equal to 1, but when breadth is grater than one the model inclusion of a minimal chain will guarantee that the resulting configuration is a chain.

Intensity as a percentage or proportion

When Intensity is expressed as a percentage or proportion, some slight modifications are needed in the previous formulation. In that case constraints in (4.20) are replaced by (4.34) and (4.35), which forces the number of float nurses that each unit can receive to range between the minimum and maximum value. Also, constraints (4.21) and (4.22) are replaced by (4.36) and (4.37) which guarantee only a number of n_j nurses may float by unit. These constraints are shown below.

$$\sum_{i\in I} w_{ij} - |I_j| \geq \alpha_j, \quad \forall j \in J \tag{4.34}$$

$$\sum_{i\in I} w_{ij} - |I_j| \leq \omega_j, \quad \forall j \in J \tag{4.35}$$

$$\sum_{\substack{j\in J\\ j\neq j'}} w_{ij} = Breadth, \quad \forall i \in I_{j'}^{n_{j'}}, \quad j' \in J \tag{4.36}$$

$$\sum_{\substack{j\in J\\ j\neq j'}} w_{ij} = 0, \quad \forall i \in I_{j'} \backslash I_{j'}^{n_{j'}}, \quad j' \in J \tag{4.37}$$

4.3.1 A heuristic to create a chain

The previous constraints to find the best chain result in a higher complexity for the NRP model, as they directly impact the problem's size and solution time. Moreover, the number of possible chain configurations increases with the number of units and intensity. Thus, besides the best chain model, we propose a heuristic based on the problem's parameter to create a chain. This heuristic works under the assumption that it is better to connect units with high and low capacities. The idea is to create a rule to use instead of, what we will call from now on, the *standard chain*, which is the chain given by default in its original order, i.e. the standard chain is at the right of Figure 4.1 where the order of units is 1-2-3-4, and by shifting that order to 4-1-3-2, we obtain a different chain. This heuristic will help to obtain a fixed chain configuration that will work with the base model. The heuristic is described as follows:

- Step 1: Get the problem information: nurses by unit ($|I_j|$), regular working hours of nurse i in the planning horizon (W_i), target demands (Q_{jdp}^{NPR}).

- Step 2: Calculate the maximum available hours per unit $UH_j = \sum_{i\in I_j} W_i$, and the total hours of target demand per unit UD_j. Note that that the demand Q_{jdp}^{NPR} is given in number of nurses for every unit, day, and period. So, parting from the logic that requiring one nurse in any hour of the day requires an hour of work of that nurse, and given than the parameter Q_{jdp}^{NPR} is in term of number nurses, we multiply it by a factor of 1 hour/nurse. In other words, if $Q_{111}^{NPR} = 3$ then the number of hours is $(3 \quad nurses) * (hours/nurse) = 3hours$. Then we could obtain the total demand in hours for the whole planning horizon, in a unit j by: $UD_j = \sum_d \sum_p ((Q_{jdp}^{NPR} nurses) * (hours/nurse))$.

– Step 3: Calculate the ratio $r_j = UH_j/UD_j$ and sort units by the calculated ratio in descending order.

– Step 4: Create links between high and low ratio pair of units, so the highest ratio unit is linked with the lowest and the second higher is linked with the second lower ratio unit. If $r_4 > r_3 > r_2 > r_1$ links are created between the units 4-1 and 3-2.

– Step 5: Create additional links between the created pairs and close the sequence with the unit of highest r_j. The chain then becomes 4-1-3-2-4, which is the chain represented in the right side of Figure 4.2. In case of an odd number of units put the last unit before closing the sequence.

For a better understanding of this heuristic procedure lets see the following example:

- *Step 1.* Consider the example data in Table 4.4. We have a problem with four units, each one of has a number of dedicated nurses. To simplify the example we assume the maximum number of regular working hours (W_i) for every nurse inside a unit is the same e.g. all nurses in unit 1 can work a maximum of 62 regular hours in the planning horizon according to the table.

Table 4.4: Example data for heuristic

Unit (J)	$\lvert I_j \rvert$	W_i	UH_j	UD_j	*Ratio* (UH_j/UD_j)
1	9	62	558	961	0.58
2	11	62	682	808	0.84
3	13	70	910	927	0.98
4	17	72	1224	842	1.45

- *Step 2.* With the data in columns 2 and 3 it is possible to calculate the total available hours per unit (UH_j) by simply multiplying the number of nurses by the working hours, which product is shown in column 3. There are target demands for nurses, previously denoted in the parameters' description of the base model as Q_{jdp}^{NPR}. Even if these demands are given as number of nurses necessary on an unit on a given day and period, in time terms a nurse needed on a given time translates in an hour of work demanded. This means that if, for example $Q_{111}^{NPR} = 4$ it is correct to say that 4 nurses are required on unit 1, day 1, period 1, and also correct to say 4 hours of work are necessary. By this logic of means, the total demanded hours can be calculated as $UD_j = \sum_d \sum_p ((Q_{jdp}^{NPR} nurses) * (hours/nurse))$ as stated before, and it for this example, we have provided the demanded hours in column 5.

- *Step 3.* In the last column of Table 4.4 the ratio of demanded hours over available hours is calculated. Ordering the ratios in column 6 the order becomes $r_4 > r_3 > r_2 > r_1$.

- *Step 4.* Links are created between units with higher and lower ratio, see the red arrows in Figure 4.2 (a) and (b) that shows links connecting 4-1 and 3-2.

- *Step 5.* Then, the link 1-3 in blue is added to connect the previous created pairs and finally the 2-4 link is also added to close the chain. Figure 4.2 (a) can help to better perceive creation of links, and (b) to perceive the final chain that was created.

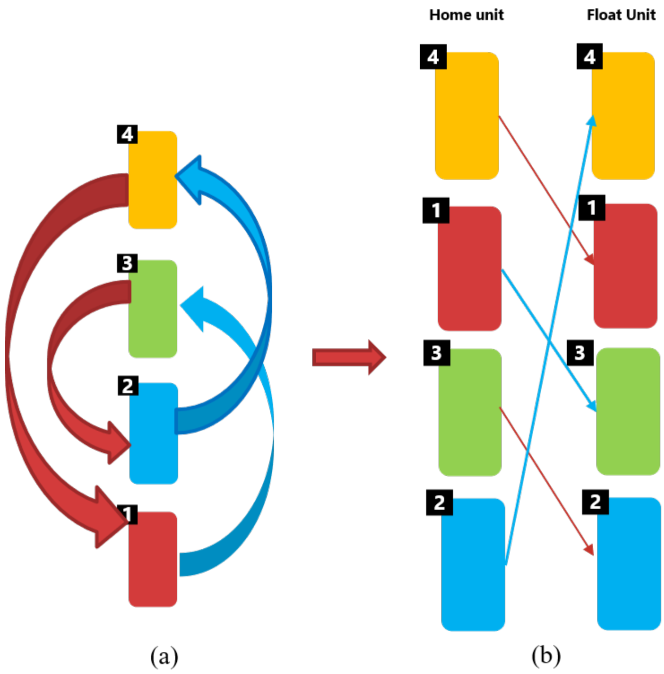

Figure 4.2: Illustration of heuristic to create a chain

The previous heuristic works to create a minimal chain when Breadth is 1. When breadth is different to 1, the order of units will be maintained and the assignments made correspondingly. So for example in Figure where Intensity=2 and breadth=2 in Figure 4.3 we see the standard chain (a) and the chain formed by the heuristic (b) and note the difference is in the order of units, and therefore in the cross-training assignments.

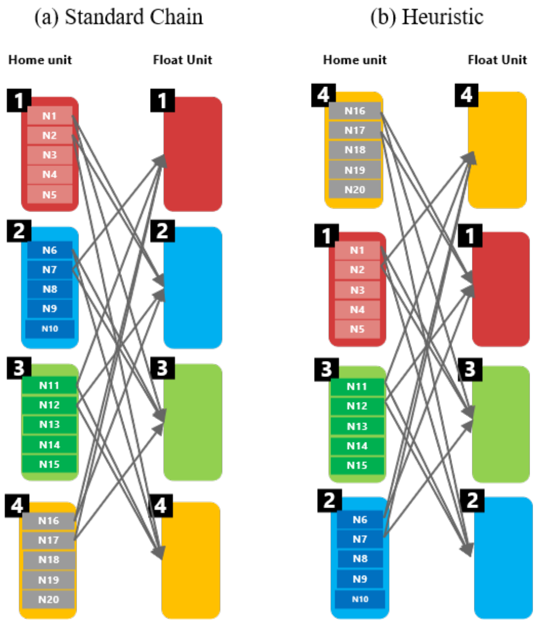

Figure 4.3: Standard Chain and heuristic for an example with Breadth = 2 and Intensity = 2

In the next section we describe the methodology developed to analyze the model variations previously proposed and the methods to generate the problem instances.

CHAPTER 5

Computational experiments

In this chapter, we describe the experiments that will help us analyze cross-training flexibility in the midterm nurse scheduling problem through the proposed model variations, CT policies, and indicators. Each experiment with its main objectives are described as follows:

1. *Preliminary experiments*: In this section, we perform experiments in two problems of different sizes; the objectives with this experiment are: a) help to understand the models by providing a better view of the variables and parameters; b) understand the relevance of the proposed model variations; c) compare the benefits of cross-training flexibility of the different CT policies; d) provide a view of the influence of parameters like intensity and H on savings, and finally; e) visualize the trade-off between savings and the investment in CT flexibility for different policies.

2. *Tests of selected flexibility policies*: After the preliminary experiments, we take the policies with the most promising results and test a new set of nine problems with a different number of units, intensities, breadth, H, and demands. The main objectives of this experiment are a) analyze the extent to which breadth, intensity, and H contribute to achieving the benefits of flexibility; b) compare the results of the chaining policies, standard chain, heuristic, and best chain model to determine if differences among them exist; c) visualize the benefits of the flexibility as a decision policy and compare them with the chaining policies in terms of the investment in cross-trainings.

3. *Test for different investment levels using the flexibility as a decision policy*: in this section, we use the model where flexibility is a decision to analyze the trade-off of the total cost and different levels of investment. We test this method in three instances from the previous experiments and comment on the minimum percentage of investment required to obtain maximum, average and minimum total costs. Finally, we analyze the proportion of time worked by dedicated nurses in every unit for each selected problem.

All experiments for the small instance were coded in Python 3.7, solved with CPLEX 12.1, and run on a computer with an Intel Xeon Six Core 5680 processor at 3.33Ghz and 36 GB RAM using one thread.

5.1 Preliminary experiments

First, we will perform experiments for two different instance sizes, that for easier reference we will call I1 and I2. Instance I1 works as an illustration and intends to provide a better view of the problem variables and parameters. Instance I2 is a large instance, closer to those of a real life setting that we will use to examine different levels of certain parameters in different policies and over which we will center our initial analysis. Breadth, understood as the number of float units a nurse is applicable to, is set to 1 for both instances. Next, we describe in detail the parameters of each instance.

Small instance (I1)

Problem I1 will be solved for nine different CT policies: (a) No flexibility, (b) chaining (Intensity=3), (c) chaining (Intensity=2), (d) reciprocal Pairs (Intensity=3), (e) n-to-all (n=3), (f) One-for-each, (g) full flexibility, and finally, (h) flexibility as a decision. Policies from (a) to (g) are forced using the base model, and in policy (h) flexibility is a decision. Each one of these policies is respectively represented in Figure 5.1. Note that the reciprocal pairs policy requires an even number of units and our example has only 3, because of this we only cross-trained nurses from two of the units as in Figure 5.1 (d).

For this example, let us consider three units, denoted by numbers 1, 2, and 3. The first unit has ten nurses, and the remaining two have five nurses each, for a total of 20 nurses. To ease the problem understanding, nurses are numbered from 1 to 20, and placed in their respective home units, as shown in the left of Figure 5.1 (a), where we can also see there are no links (in grey in other policies), as no flexibility is considered. Nurses are to be scheduled for four days, where each day consists of seven periods, and four shifts are possible. Given the information above we have the following sets:

Sets
$I = \{1, 2, ..., 20\}$ Set of nurses
$S = \{1, 2, 3, 4\}$ Set of shifts
$J = \{1, 2, 3\}$ Set of units
$D = \{1, 2, 3, 4\}$ Set of days
$P = \{1, 2, 3, 4, 5, 6, 7\}$ Set of periods

Next, we provide the parameters for the problem. Table 5.1 shows scalar parameters, and Tables 5.2 and 5.3 parameters related to shifts. The minimum and target demands, respectively represented by Q_{jdp}^{min} and Q_{jdp}^{NPR}, are tailor-made for this instance (as it was easier to verify the constraints that way). The reader might note the demands are repeated for all units, something made purposely to ease the understanding of this preliminary experiment. So, given the expected demand values E_{jdp} demands are generated as follows:

$$Q_{jdp}^{min} = E_{jdp} \qquad \forall j \in J, d \in D, p \in P \tag{5.1}$$

$$Q_{jdp}^{NPR} = Q_{jdp}^{min} + 1 \qquad \forall j \in J, d \in D, p \in P \tag{5.2}$$

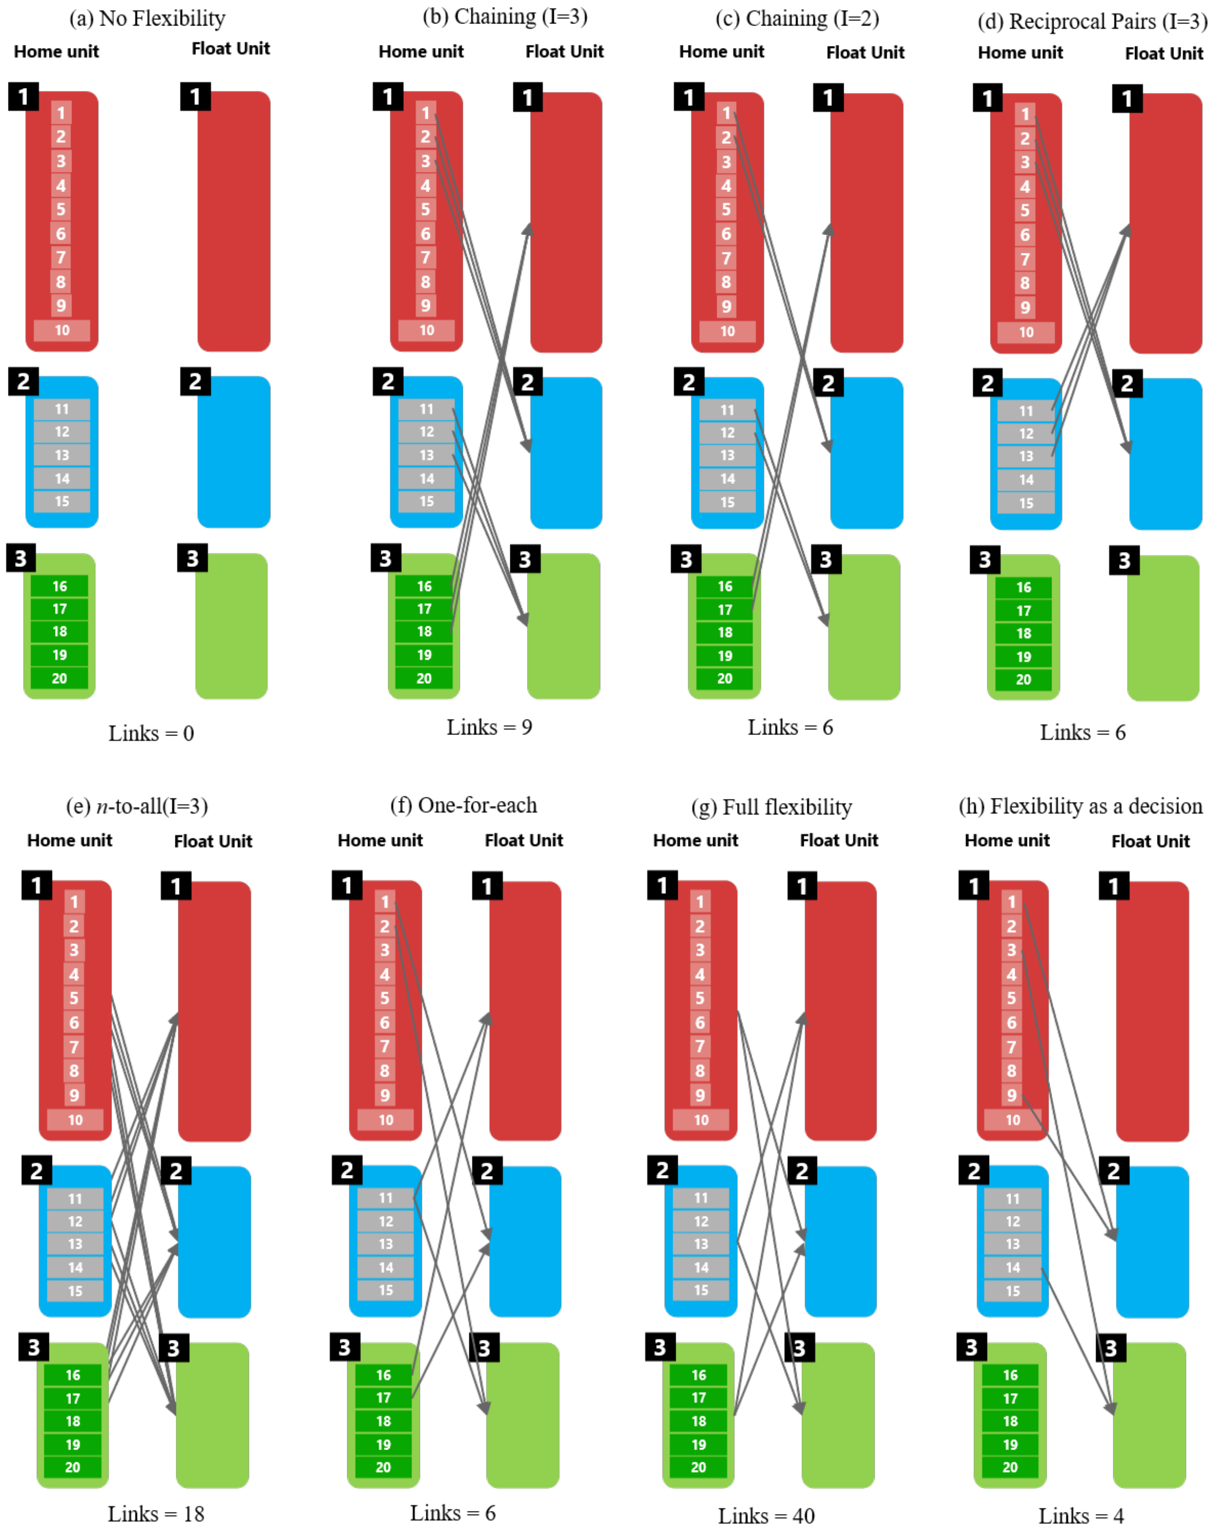

Figure 5.1: Illustration of cross-training policies for Instance I1

Table 5.1: Scalar parameters

Description	Parameter	Value
Minimum rest time between shifts (hours)	R_{min}	1
Minimum proportion of working time in Home unit	H	0.1
Costs of each nurse missing to cover minimum demand ($)	$C^{\Delta min}$	3
Costs of each nurse missing to satisfy target demand ($)	$C^{\Delta NPR}$	1
Costs for each overtime hour ($)	C^{ot}	2
Minimum number of consecutive days a nurse is assigned a unit	$Stay^{min}$	1
Maximum stretch of consecutive on duty working days	Str^{max}	3
Regular working hours of nurse i in the planning horizon	W_i	10
Maximum overtime hours of nurse i in the planning horizon	OT_i^{max}	10
1, if nurse i works shift s in unit j on the last day of the previous planning horizon, 0 otherwise	y_{isj0}	0

Table 5.2: A_{sp} parameter

		Period						
		1	*2*	*3*	*4*	*5*	*6*	*7*
Shift	*1*	1	1	0	0	0	0	0
	2	0	0	1	1	0	0	0
	3	0	0	0	0	1	1	0
	4	0	0	0	0	0	0	1

Table 5.3: Shifts specifications

$Shift$	Beg_s	End_s	L_s
1	1	2	2
2	3	4	2
3	5	6	2
4	7	7	1

Table 5.4: Minimum and target demands

day	period	Unit 1		Unit 2		Unit 3	
		Q_{jdp}^{min}	Q_{jdp}^{NPR}	Q_{jdp}^{min}	Q_{jdp}^{NPR}	Q_{jdp}^{min}	Q_{jdp}^{NPR}
1	1	1	2	1	2	1	2
1	2	1	2	1	2	1	2
1	3	2	3	2	3	2	3
1	4	2	3	2	3	2	3
1	5	1	2	1	2	1	2
1	6	1	2	1	2	1	2
1	7	1	2	1	2	1	2
2	1	2	3	2	3	2	3
2	2	2	3	2	3	2	3
2	3	1	2	1	2	1	2
2	4	1	2	1	2	1	2
2	5	1	2	1	2	1	2
2	6	1	2	1	2	1	2
2	7	1	2	1	2	1	2
3	1	0	1	0	1	0	1
3	2	0	1	0	1	0	1
3	3	1	2	1	2	1	2
3	4	1	2	1	2	1	2
3	5	2	3	2	3	2	3
3	6	2	3	2	3	2	3
3	7	1	2	1	2	1	2
4	1	1	2	1	2	1	2
4	2	1	2	1	2	1	2
4	3	1	2	1	2	1	2
4	4	1	2	1	2	1	2
4	5	1	2	1	2	1	2
4	6	1	2	1	2	1	2
4	7	1	2	1	2	1	2

The time limit for this experiment is set to 2 minutes.

Large instance (I2)

Problem I2 will be solved for nine different cross-training policies, three levels of intensity and two levels of H, which is the minimum proportion of regular working time a nurse has to work in its home unit (continuity of care). An overview of the 35 scenarios is shown in Table 5.8. This problem has similar parameters to the one proposed in the article of Fügener et al. (2018), with the difference that 50 nurses are considered instead, and the number of nurses in each unit is different. So, there are 50 nurses in total, four units, nine shifts and 15 days with 24 periods each. The full set of parameters of the problem is shown from Tables 5.5 to 5.7. Note that it is possible to obtain A_{sp} from the information on Table 5.7 which is why it was not included any table for this parameter.

Table 5.5: Nursing personnel classified by home units

Unit (j)	Number of nurses with $HU_i = j$	W_i (in hours)
1	9	62
2	11	62
3	13	70
4	17	72

Table 5.6: Scalar parameters

Description	**Parameter**	**Value**
Minimum rest time between shifts (in hours)	R_{min}	11
Costs of each nurse missing to cover minimum demand ($)	$C^{\Delta min}$	1000
Costs of each nurse missing to satisfy target demand ($)	$C^{\Delta NPR}$	10
Costs for each overtime hour ($)	C^{ot}	100
Minimum number of consecutive days a nurse is assigned a unit	$Stay^{min}$	2
Maximum stretch of consecutive on duty working days	Str^{max}	6
Maximum overtime hours of nurse i in the planning horizon	OT_i^{max}	6
1, if nurse i works shift s in unit j on the last day of the previous planning horizon, 0 otherwise	y_{isj0}	0

Table 5.7: Shift Parameters

Shift	Beg_s	End_s	L_s
1	1	8	8
2	1	6	6
3	3	6	4
4	3	8	6
5	7	11	5
6	7	14	8
7	9	16	8
8	11	16	6
9	17	24	8

Table 5.8: Overview of scenarios

H	Intensity	Chaining	Reciprocal Pairs	n-to-all	Full Flexibility	One for each	Best Chain	Heuristic	Flexibility as a decision	No CT
0.5	1	*Scenario 1*	*Scenario 4*	*Scenario 7*			*Scenario 12*	*Scenario 30*		
0.5	3	*Scenario 2*	*Scenario 5*	*Scenario 8*		*Scenario 11*	*Scenario 13*	*Scenario 31*		
0.5	50%	*Scenario 3*	*Scenario 6*	*Scenario 9*	*Scenario 10*		*Scenario 14*	*Scenario 32*	*Scenario 15*	*Scenario 16*
0.8	1	*Scenario 17*	*Scenario 20*	*Scenario 23*			*Scenario 27*	*Scenario 33*		
0.8	3	*Scenario 18*	*Scenario 21*	*Scenario 24*		*Scenario 26*	*Scenario 28*	*Scenario 34*		
0.8	50%	*Scenario 19*	*Scenario 22*	*Scenario 25*			*Scenario 29*	*Scenario 35*		

The minimum and target demands, respectively represented by Q_{jdp}^{min} and Q_{jdp}^{NPR}, were randomly generated and fixed for this example. As in the base model equations 4.13 to 4.15 were used to generate demands and the parameter u_{jd} was considered with $n = 2$, this is, $u_{jd} \sim U_{jd}\{-2; 2\}$. The breadth in all cases, with exception of the n-to-all policy, will be set to 1. When Intensity = 50% it means half of the nurses in each are cross-trained to work in other units. The number of cross-trained nurses in each unit (n_j) will be calculated as $n_j = \lceil Intensity * |I_j| \rceil$. So for this example $n_1 = 5, n_2 = 6, n_3 = 7, n_4 = 9$, for a total of 27 cross-trained nurses. Finally, the time limit was set to four hours in all scenarios.

Results discussions on preliminary experiments

First, lets comment on the results of instance I1. The objective function value for the tested policies, and the number of cross-trainings is shown in Table 5.9. The number of cross-trainings, with exception of policy (h) is easy to calculate according to the intensity and number of units. As it can be noted in Table 5.9 it was possible to obtain better objective values by applying flexibility with cross-training. Policies (b), (e), and (h) obtained all the possible benefits of full flexibility (g). It was policy (h) the one that required less links to obtain all those benefits with only four links for the greatest possible benefit i.e. by training only four nurses to work in another units it is possible to obtain the same benefit of 40 trained nurses required for full flexibility.

Table 5.9: Results for instance I1

CT Policy	Objective	Num CT
(a) No flexibility	78	0
(b) Chaining (I=3)	72	9
(c) Chaining (I=2)	74	6
(d) Reciprocal Pairs (I=3)	75	6
(e) n-to-all (I=3)	72	18
(f) One-for-each (I=3)	74	6
(g) Full flexibility	72	40
(h) Flexibility as decision	72	4

In Table 5.10 we can see the solution schedule for CT policy (a), for each unit. The numbers between braces indicate nurses that cover the given shift on the indicated unit and day. A look into Table 5.10, allows to perceive how in this case, there is not problem to supply the minimum demand in unit 1, unlike in unit 2 and 3 in period 7 (in this case covered in shift 4) for the first three days, where as shown in their respective tables no nurses could be assigned. In unit 1 it was not possible to reach the target demand in the last period in any of the days, and as it was expected because the minimum could not be reached, in unit 2 and 3 it was not possible to reach the target in any day and period by 1 nurse. Table 5.11 shows a summary per unit of the Δ_{jdp}^{min}, Δ_{jdp}^{NPR} values that the reader can use to verify the objectives and assignments.

Table 5.10: Schedules for policy (a), I1

		Day			
		1	**2**	**3**	**4**
Unit 1	**1**	[5, 6]	[4, 8, 9]	[10]	[6, 9]
	2	[2, 4, 8]	[2, 5]	[6, 9]	[7, 8]
	3	[7, 10]	[3, 10]	[1, 4, 7]	[3, 5]
	4	[1]	[1]	[3]	[2]
Unit 2	**1**	[12]	[12, 14]	-	[15]
	2	[11, 13]	[13]	[11]	[13]
	3	[15]	[15]	[12, 14]	[11]
	4	-	-	-	[14]
Unit 3	**1**	[17]	[17, 19]	-	[20]
	2	[16, 18]	[18]	[16]	[18]
	3	[20]	[20]	[17, 19]	[16]
	4	-	-	-	[19]

Table 5.11: Summary of Δ_{jdp}^{min} and Δ_{jdp}^{NPR}

	Unit 1	**Unit 2**	**Unit 3**
Δ_{jdp}^{min}	0	3	3
Δ_{jdp}^{NPR}	4	28	28

Table 5.12: Schedules for policy (b), I1

		Day			
		1	**2**	**3**	**4**
Unit 1	**1**	[6, 7]	[9, 10]	-	[9, 10]
	2	[5, 9, 10]	[3, 4]	[1]	[5]
	3	[2]	[8]	[4, 5, 7]	[4, 7]
	4	[8]	[6]	[8]	[6]
Unit 2	**1**	[14]	[12, 13]	-	[3]
	2	[13, 15]	[1]	[14]	[2]
	3	[1]	[11]	[2, 15]	[15]
	4	[3]	[14]	[13]	[11]
Unit 3	**1**	[11]	[17, 20]	-	[16]
	2	[17, 18]	[16]	[12]	[19]
	3	[19]	[19]	[17, 20]	[12]
	4	[16]	[18]	[18]	[20]

Table 5.13: Schedules for policy (h), I1

		Day			
		1	**2**	**3**	**4**
Unit 1	**1**	[1, 4]	[5, 8, 10]	-	[4, 10]
	2	[6, 10]	[4]	[7, 8]	[8]
	3	[2]	[7, 9]	[5, 6]	[3, 5]
	4	[7]	[2]	[2]	[6]
Unit 2	**1**	[13]	[11, 12]	-	[14]
	2	[9, 11]	[13]	[14]	[11]
	3	[12]	[1]	[13, 15]	[9]
	4	[15]	[15]	[1]	[12]
Unit 3	**1**	[18]	[17, 19]	-	[17]
	2	[14, 16]	[20]	[16]	[18]
	3	[19]	[3]	[3, 17]	[19]
	4	[20]	[18]	[20]	[16]

Now, let us take a look into the schedules for policy (b) in Table 5.12. As observed, there is no problem covering the minimum demand in any case, as some nurses from unit 1 are now floated to unit 2 (see nurses 1,2 and 3) and from unit 2 to unit 3 (see nurse 11). Here we can see that the objective value is 72, as the full flexibility case, and that effectively four nurses were scheduled to other units, even if 9 in total were cross-trained with the policy. This result is coherent with policy (h), where only four nurses were cross-trained and effectively scheduled to other units; see nurses 1, 3, 9, 14 in Table 5.13. The main take from the results in instance I1, consistently with the literature, is that it is not necessary to add total flexibility to obtain most or close to all of its benefits. In this case, cross-training four nurses to other units is enough to obtain all the possible benefits.

Results for scenarios of I2 are presented in Table 5.14. This table shows the objective value, gap percentage, running time, number of cross-trainings (links), savings percentage (higher values in bold), and savings by the number of links for each scenario. The savings percentage of a scenario i was calculated as $\left(\frac{O_{16}-O_i}{O_{16}}\right)*100\%$, where O_{16} is the objective for scenario 16 (No flexibility) and O_i is the objective of scenario i. The number of links can be obtained as the product of the intensity, breadth and number of units $Intensity*Breadth*|J|$. For example in Scenario 1, when Intensity is 1, according to the previous equation, the number of links is 1*1*4 = 4 (remember that breadth is fixed to 1 for these experiments). For scenarios when the intensity is a percentage as in Scenario 3 this number of links is obtained as $(\sum_j n_j) * Breadth$, when the first term corresponds to the total cross-trained nurses, calculated as in the previous subsection. The savings by the number of links is obtained by dividing the previous calculated savings percentage over the number of cross-trainings (links).

A general view on the results lets to perceive higher costs when there is no cross-training and a lower cost in scenario 15, with flexibility as a decision. In the latter scenario, we added constraints so that it was possible to invest in flexibility. For this run, as there is no decision-maker to indicate a possible value for the number of desired cross-trainings, we included constraint (4.18) in the objective function looking to find the solution that minimized both the costs and the number of links in scenario 15. The overall cost was 2450 using 22 links with a gap of 16% in four hours, which was the best objective value among all scenarios, including scenario 16 (full flexibility), where the costs were 2910 with a gap of 29% in the allowed time. Most of the scenarios were optimally solved in the given time, but some significant gaps stand out in the full flexibility, flexibility as a decision, and one of the n-to-all scenarios. Furthermore, gaps above 5% are present in scenarios where the intensity is 50% (half of the nurses in each unit are cross-trained) and $H = 0.5$ due to a more significant number of links; see scenarios 3, 9, 14, and 30. When looking at solution times, we see shorter times were associated with low intensities in most policies (an exception happened in scenario 5).

Table 5.14: Results for instance I2

Policy	Intensity	H	Scenario	Objective	Gap (%)	Time (s)	Links	Savings (%)	Savings (%) /# links
No CT	0	-	16	20350	2.80	14400	0	0.00	-
Full flexibility	0	0.5	10	2910	29.47	14400	150	**85.70**	0.57
Flexibility as decision	-	0.5	15	2450	16.23	14400	22	87.96	3.99
Standard Chain	1	0.5	1	10360	0.00	766	4	49.09	12.27
	3	0.5	2	4600	0.00	1800	12	77.40	6.45
	50%	0.5	3	2670	7.91	14400	27	**86.88**	3.22
	1	0.8	17	13370	0.00	743	4	34.30	8.57
	3	0.8	18	9580	0.00	2334	12	52.92	4.41
	50%	0.8	19	4420	0.00	7674	27	78.28	2.90
Reciprocal Pairs	1	0.5	4	11070	0.00	2487	4	45.60	11.40
	3	0.5	5	5470	1.65	14400	12	73.12	6.09
	50%	0.5	6	5380	0.00	2185	27	73.56	2.72
	1	0.8	20	13570	0.00	2179	4	33.32	8.33
	3	0.8	21	10250	0.00	5575	12	49.63	4.14
	50%	0.8	22	6910	0.00	5438	27	66.04	2.45
n-to-all	1	0.5	7	5280	0.00	8314	12	74.05	6.17
	3	0.5	8	4520	0.00	11385	36	77.79	2.16
	50%	0.5	9	2740	11.04	14400	81	**86.54**	1.07
	1	0.8	23	10020	0.00	1794	12	50.76	4.23
	3	0.8	24	5220	0.19	14400	36	74.35	2.07
	50%	0.8	25	4380	1.14	14400	81	78.48	0.97
One-for-each	3	0.5	11	4610	0.59	14400	12	77.35	6.45
	3	0.8	26	9750	0.00	1635	12	52.09	4.34
Best Chain	1	0.5	12	10340	0.00	14400	4	49.19	12.30
	3	0.5	13	4640	2.40	14400	12	77.20	6.43
	50%	0.5	14	2630	7.25	14400	27	**87.08**	3.23
	1	0.8	27	13350	0.00	3169	4	34.40	8.60
	3	0.8	28	9650	2.98	14400	12	52.58	4.38
	50%	0.8	29	4490	3.56	14400	27	77.94	2.89
Heuristic	1	0.5	31	10340	0.00	1341	4	49.19	12.30
	3	0.5	32	4580	0.00	11175	12	77.49	6.46
	50%	0.5	30	2590	5.47	14400	27	**87.27**	3.23
	1	0.8	34	13350	0.00	747	4	34.40	8.60
	3	0.8	35	9560	0.00	4952	12	53.02	4.42
	50%	0.8	33	4410	0.45	14400	27	78.33	2.90

Looking at the savings (%) vs the number of cross-trainings (Links) plotted in Figure 5.2, it can be observed that it is possible to achieve almost all the benefits of full flexibility with less number of cross-trainings, in this case with policies like n-to-all, chaining (standard chain, heuristic, best chain) and flexibility as a decision. It can also be noted that scenario 15, where flexibility is a decision, obtained more savings than the full flexibility case (Scenario 10), which can be explained by gaps superior to 15% in each of those policies, which also suggests higher savings are possible.

To better look at the incidence of the minimum proportion of working time in the home unit (H) over the results, Figure 5.3 shows all the policies in both proposed levels. The blue dots' position above the red marks suggests that the greater the proportion of time nurses spent in a home unit, the less are the savings. Table 5.15 shows the average savings by H level for every policy and the differences between them. In the table, we corroborate what was suggested in the plot, as in

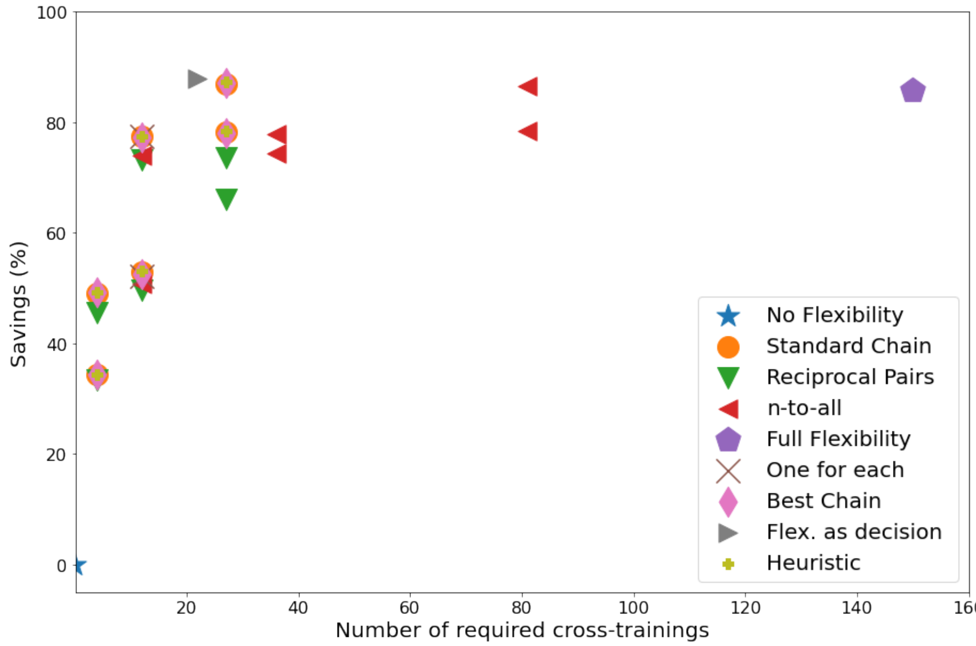

Figure 5.2: Savings (%) vs. Number of required cross-trainings by policy

every CT policy the largest average savings were when $H = 0.5$, making the average savings for all policies 72.1% when $H = 0.5$ and 56.1% when $H = 0.8$. Based on the difference column, we could say that policies with lower H yielded 16.1% more savings on average. The average column of the same table allows us to see that the policy with more average savings was the n-to-all, followed by one-for-each and the heuristic, without considering the number of links used.

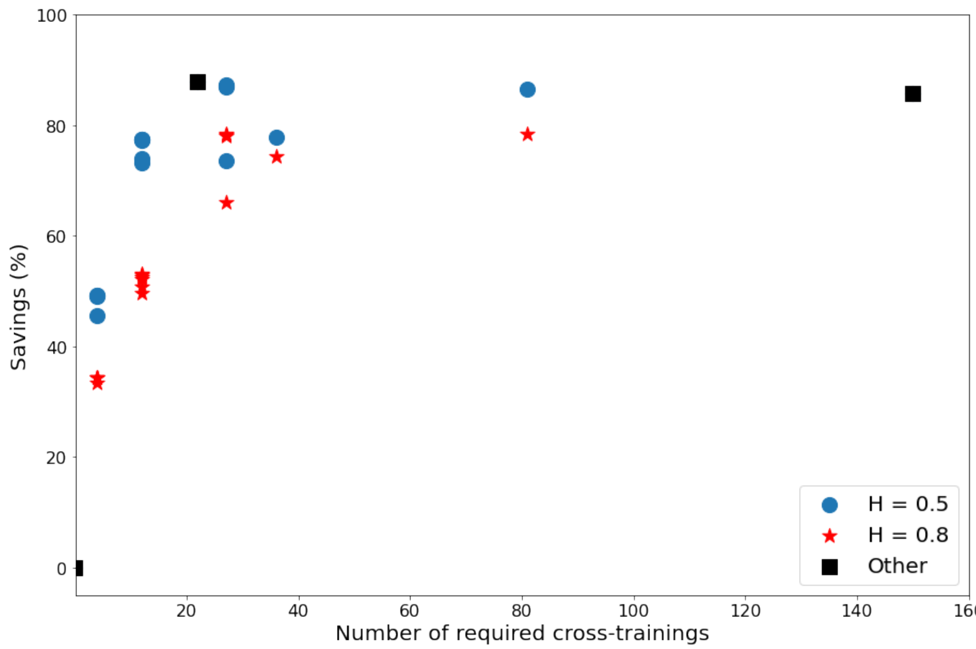

Figure 5.3: Savings (%) vs. Number of required cross-trainings classified by H

When looking at the level of intensity, greater intensities translated into greater savings in each policy. The savings, averaging all policies, were 45%, 66%, and 77% for intensities of 1, 3, and 50%, respectively. Figure 5.4 shows the different intensities in red, green, and blue. In the figure, all the blue markers (Intensity = 50%) go over 60% and some of them have the most significant possible

Table 5.15: Average savings by levels of H and policies

Policy	H **0.5**	H **0.8**	**Difference**	**Average**
Standard Chain	71.1%	55.2%	16.0%	63.1%
Reciprocal pairs	64.1%	49.7%	14.4%	56.9%
n-to-all	79.5%	67.9%	11.6%	73.7%
One for each	77.3%	52.1%	25.3%	64.7%
Best Chain	71.2%	55.0%	16.2%	63.1%
Heuristic	71.3%	55.2%	16.1%	63.3%
Average	72.1%	56.1%	16.1%	64.1%

savings, which is consistent with the greater average savings corresponding to intensities of 50%.

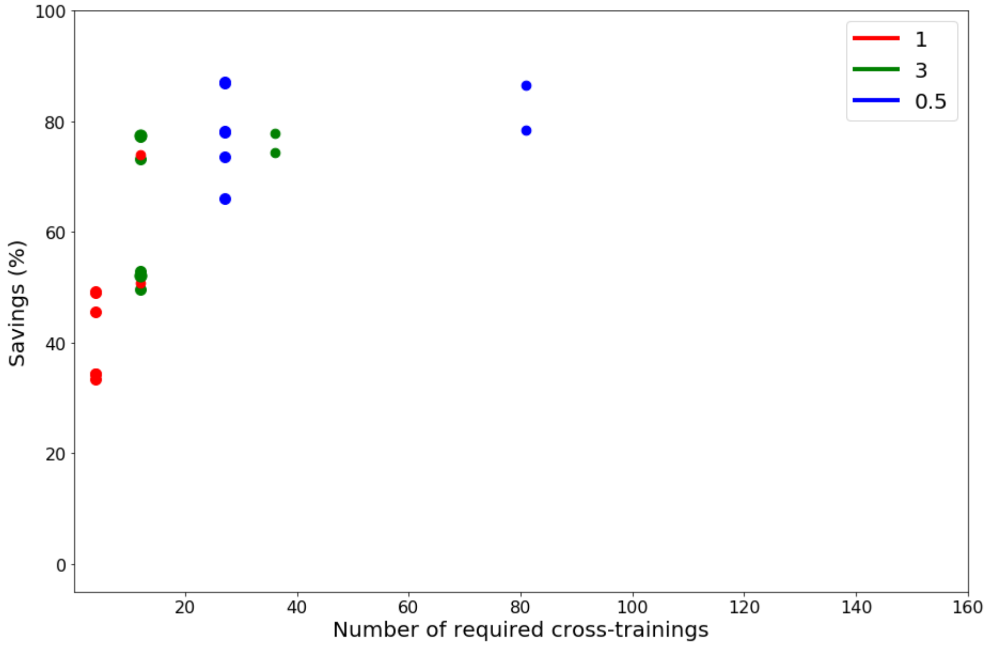

Figure 5.4: Savings (%) vs. Number of required cross-trainings by Intensity

By combinations of H and intensity, some policies performed better than others. However, to compare them only in terms of savings would not be fair as the number of links differs and is greater for some policies. So, the percentage of savings that each link represents for every result could be a better measure, shown in the last column of the results table. Table 5.16 shows the best policy for each combination of H and Intensity considering savings (%) and savings (%)/$\#links$. According to the table, the n-to-all performs better in five of the six combinations if only savings are considered. However, because the n-to-all policy uses a large number of cross-trainings, when calculating the savings that each link represents, the policies that use chaining and the one-for-each appear as the best policies. If we compare the policies that use the same number of links (4, 12, and 27 links), a look into the results table shows that the policies using chains performed slightly better for those numbers of links.

Table 5.16: Best policy by combination of H and intensity

H	**Intensity**	**Best policy**	
		By savings	**By savings/# of links**
0.5	1	n-to-all	Best chain, Heuristic
0.5	3	n-to-all	Heuristic, Standard Chain, One-for-each
0.5	50%	Heuristic	Heuristic, Standard Chain, Best Chain
0.8	1	n-to-all	Best chain, heuristic
0.8	3	n-to-all	Standard Chain, heuristic
0.8	50%	n-to-all	Heuristic, Standard Chain

Analysis of proposed variations

Next, we analyze the proposed model variations: flexibility as a decision, best chain, and the heuristic. As previously mentioned, the flexibility as a decision model in scenario 15 showed the greatest savings among the other policies with 87.96% using 22 cross-trainings, despite having a gap of a little over 16%. Chapter 4 proposed a heuristic that created a chain, not assigning links between units in a standard order but using the demanded hours for each unit and nurses' availability. When comparing this approach with the standard chain and the best chain in Table 5.17 (best values in bold), we note slight improvements of the heuristic results over the standard chain in every case, and in some cases over the best chain. The best chain model results should be the best among the three policies, but higher gaps were obtained and only got to equal the heuristic in two H and intensity combinations. Finally, important differences in the solution times were found, with respective average values of 4619, 12528, and 7835 seconds for the standard chain, best chain, and the heuristic.

Table 5.17: Chaining policies comparison

H	**Intensity**	**Standard Chain**	**Best Chain**	**Heuristic**
0.5	1	10360	**10340**	**10340**
0.5	3	4600	4640	**4580**
0.5	50%	2670	2630	**2590**
0.8	1	13370	**13350**	**13350**
0.8	3	9580	9650	**9560**
0.8	50%	4420	4490	**4410**

Finally, when looking at the scenario with the best results in terms of savings and number of cross-training links as objectives, four scenarios were better than all the others but not better (or worse) among them. These scenarios can be found in Figure 5.5 highlighted by the blue markers. Note as well, by the overlapping over different shapes of markers, that differences were relatively small. This information can be helpful to a decision-maker who wants to analyze the trade-off of cross-training and savings of all policies for this instance.

From the previous experiment we can summarize the following observations:

- The parameter H has an important impact on the benefits of flexibility.
- The chaining policies, namely Best chain, Heuristic and Standard chain showed to have the best results for the different combinations of H and intensity.

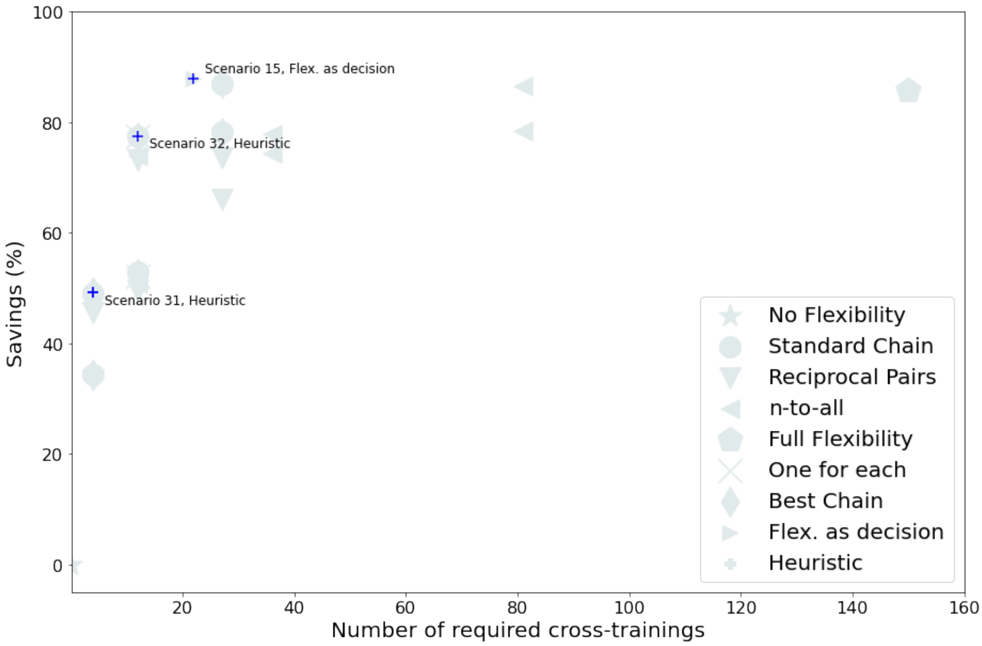

Figure 5.5: Non-dominated scenarios

- The flexibility as a decision model was capable to find the greatest percentage of savings among the analyzed policies while minimizing the number of cross-trainings. Despite this, the higher complexity of this extension was reflected in a gap a little over 16%.

5.2 Experiments of selected cross-training flexibility policies

From now on when talking about a problem we do not refer to our research problem (i.e mid-term nurse scheduling) but to a setup of fixed parameters with a configuration of nurses distributed in units. When referring to an instance of a problem we refer to the previous setup combined with a level of Intensity, Breadth, H, policy and demand.

Based on the results from the previous experiment, we decided to select the best policies to extend our analysis to other problems. To make the selection of policies, we considered the savings percentage and the number of cross-trainings in each policy. In the previous section, we noted it would not be fair to judge the results only in the light of the total savings, as the number of cross-trainings might differ by policy. For the previous reason, we added the savings (%)/#links column to represent the percentage that every training contributed to the total savings. As seen in Table 5.16, the chaining policies (Best chain, standard chain and heuristic) showed to provide better results. For the reasons mentioned above we selected them together with the flexibility as a decision policy to perform experiments in this section. We decided not to test the one-for-each policy as the number of cross-trainings is fixed according to the number of units, and necessarily the minimum number of nurses by unit has to be $|J| - 1$. The previous is a limitation as the problems used in this section do not have a homogeneous number of nurses among units, and this minimum number is not guaranteed.

In this experiment, we will deepen the conclusions from the preliminary experiments through nine different problems, each one of which we will test under different levels of intensity, breadth, continuity of care (H), policies, and demands. The full experiment by level of the previous parameters is described in Table 5.18. There are nine problems in total, three for each 20, 30, and 40 number of nurses as indicated at the beginning of each problem's name, i.e., problems starting by 20 have that number of available nurses. The number of units across problems can vary, as well as the distribution of nurses among those units. Levels of CT indicators like breadth and intensity respectively depend on those factors. For example, the maximum level of breadth is $|J| - 1$, where J is the set of units in a problem. On the other hand, the intensity levels in the table are set to the minimum number of nurses among units, meaning that if there are four units in the problem and the unit with fewer nurses has two, the maximum intensity level is set to two.

Table 5.18: Experiment design for selected flexibility policies

Problem	**#Units**	***Levels***					***Total Runs***
		Intensity	Breadth	H	Policies	Demands	
20I1	4	4	3	2	4	100	9600
20I2	5	2	4	2	4	100	6400
20I3	5	3	4	2	4	100	9600
30I1	4	5	3	2	4	100	12000
30I2	5	3	4	2	4	100	9600
30I3	6	3	5	2	4	100	12000
40I1	4	8	3	2	4	100	19200
40I2	5	6	4	2	4	100	19200
40I3	6	6	5	2	4	100	24000

The levels of intensity and breadth are consecutive, so if there are four levels of intensity, it means experiments will run for the cases where one, two, three and four nurses are cross-trained by unit. If the breadth level is two, then those nurses will have to be cross-trained to one and two additional units for each of the intensity numbers of nurses. We consider two levels of continuity of care like in instance I2 in the preliminary experiments (0.5 and 0.8), four policies including standard chain, heuristic, best chain and flexibility as a decision. The total runs by problem are the product of intensity, breadth, H, policies, and demands. So for example for problem 20I1 the total runs are $4 * 3 * 2 * 4 * 100 = 9600$ runs.

When the model of flexibility as a decision is used the right side of constraint (4.18) that limits the investment will be set to $Intensity * Breadth * |J|$ and constraint (5.3) will limit the breadth of every nurse. We do this to have a fair comparison among policies.

$$\sum_{j \in J} w_{ij} \leq Breadth + 1, \forall i \in I \tag{5.3}$$

A total of 100 different demands were also generated for each problem. The problems share some parameters in common as shown in Table 5.19. The parameters $Stay^{min}$ and Str^{max} can be different among problems, with minimum and maximum values of 1 and 2 for the former, and 5 and 6 for the latter. The distribution of nurses across units for every problem is shown in Table 5.20.

Demands were generated in a similar way as in the article of Fügener et al. (2018). First, as we do not have real-life data, we generated the expected demands E_{jdp} for each problem by using a discrete uniform distribution. For each problem, the expected demands were generated as in table 5.21. Note that as we are dealing with a number of nurses demanded, only the integer part of the

Table 5.19: Commom parameters for all problems

Description	**Parameter**	**Value**
Days (Planning horizon)	$\|D\|$	10
Periods	$\|P\|$	24
Shifts	$\|S\|$	5
Minimum rest time between shifts (hours)	R_{min}	8
Costs of each nurse missing to cover minimum demand	$C^{\Delta min}$	1000
Costs of each nurse missing to satisfy target demand	$C^{\Delta NPR}$	10
Costs for each overtime hour	C^{ot}	100
Regular working hours of nurse i in the planning horizon	Wi	80
Maximum overtime hours of nurse i in the planning horizon	OT_i^{max}	10
1, if nurse i works shift s in unit j on the last day of the previous planning horizon, 0 otherwise	y_{isj0}	0

Table 5.20: Distribution of nurses by unit for every problem

Problem	**Number of nurses**					
	Unit 1	Unit 2	Unit 3	Unit 4	Unit 5	Unit 6
20I1	6	6	4	4	-	-
20I2	5	4	2	5	4	-
20I3	5	5	3	3	4	-
30I1	10	8	5	7	-	-
30I2	9	6	3	5	7	-
30I3	5	6	7	3	4	5
40I1	11	11	10	8	-	-
40I2	10	7	9	6	8	-
40I3	7	6	6	6	7	8

generated numbers is considered, leaving the expected number of nurses with a pretty limited set of possibilities. So, for example, if $E_{jdp} \sim E_{jdp}\{1, 3\}$ a minimum of 1 and a maximum of 2 nurses can be expected in every hour of the day.

Table 5.21: Uniform distribution minimum and maximum by problem

Problem	**Min**	**Max**
20I1	1	3
20I2	1	3
20I3	1	3
30I1	1	4
30I2	1	3
30I3	1	4
40I1	1	4
40I2	1	3
40I3	1	3

Now, just like the authors of the base model, to simulate the effect of stochastic demand, the deterministic demand is adjusted by unit by day with a u_{jd} parameter, which is a number from a discrete uniform distribution $u_{jd} \sim U_{jd}\{-2; 2\}$. During the night shift, no deviation from the expected demand is considered, and the target demand is only one nurse over the minimum demand. The E_{jdp} values are fixed by problem, and to generate each of the 100 demands, we generate new values of u_{jd} parameter. When the resulting number of nurses is negative, the number of nurses is set to zero. Finally, the demands are generated according to expressions (5.4) to (5.6).

$$Q_{jdp}^{min} = E_{jdp} + u_{jd} \qquad \forall j \in J, d \in D, \{p \in P \mid p \leq 16\} \tag{5.4}$$

$$Q_{jdp}^{min} = E_{jdp} \qquad \forall j \in J, d \in D, \{p \in P \mid p \geq 17\} \tag{5.5}$$

$$Q_{jdp}^{NPR} = Q_{jdp}^{min} + 1 \tag{5.6}$$

In addition to the experiments previously presented, for each problem, level of H, and demand, full and no flexibility policies will also be performed. Those policies will help us to have a better view of the benefits gained by the investment in flexibility.

Results discussion for selected policies

To analyze intensity, breadth, and H, first we calculate the average benefits of the 100 demands for each problem and H level. This average benefit is calculated as the difference in average costs between the full and no flexibility policies for each problem. These average benefits are shown in 5.22. It stands out that there is no difference between full and no flexibility when $H = 0.8$ in three of the proposed problems. More specifically, considering this level of H it is not possible to grasp any benefits from adding flexibility in problems 20I3, 30I1, and 40I2. In the remaining instances, there is a reduction in the possible benefits when going from $H = 0.5$ to $H = 0.8$; this reduction is less than 5% in problems 20I1 and 40I1, 25% on average in problems 20I2, 30I2, 30I3 and 40I3, and 100% (the benefit reduced to 0) in the remaining problems. This is likely to happen as a greater H causes cross-trained nurses to have less allowed time (20 % of regular time) to work in other units thus, increasing the full flexibility cost and therefore the benefit.

Table 5.22: Average costs by problem and H level, with full and no flexibility

Problem	Avg. Cost H = 0.5			Avg. Cost H = 0.8		
	Full Flexibility	No Flexibility	Benefits	Full Flexibility	No Flexibility	Benefits
20I1	64277.8	124435.4	60157.6	64298.1	124435.6	60137.5
20I2	507245.4	571745.2	64499.8	513838.8	571745.2	57906.4
20I3	420287.5	465989.4	45701.9	466351.2	466351.2	0
30I1	258705.1	319045.1	60340	320113	320113	0
30I2	102537.4	256273.9	153736.5	145465.6	256273.8	110808.2
30I3	327411.5	444848.5	117437	355402.3	444848.5	89446.2
40I1	25059.2	62245.9	37186.7	26814.9	62245.9	35431
40I2	5162.6	50188.4	45025.8	50430.3	50430.3	0
40I3	47757.4	115099.8	67342.4	56004.4	115100	59095.6

With the values presented in Table 5.22 and the results from each instance of the experiment for each problem, we can calculate which average percentage of benefits are achieved when using a combination of parameters. For example, lets suppose that the averaged cost (for the 100 demands) when $Intensity = 1$, $Breadth = 1$ and $H = 0.5$ in problem 20I1 using the best chain was 84038.9. With this average cost and the values in Table 5.22 we can calculate the average percentage of benefits as follows:

$$\%Avg.Benefits = \left(\frac{124435.4 - 84038.9}{60157.6}\right) * 100\% = 67.15\%$$

Another way of seeing this benefit is as the average reduction of costs from the worst (no flexibility) to the best (full flexibility) case scenario. So, we could say that in the previous example, it is possible to achieve on average 67.15% of the benefits of full flexibility by using the best chain and combining the given H and CT parameters. Now that we know how to calculate the percentage of benefits we proceed to show the benefits for the different levels of parameters using Table 5.23 for $H = 0.5$ and Table 5.24 for $H = 0.8$. We see the corresponding problem and different intensity levels in rows and policies with breadth levels in columns in each table. The percentage of benefits achieved with each combination of parameters is shown at the intersection of each row and column.

We used a color scheme with tones from red to green as a visual aid to dimension the percentage of average benefits when moving along different intensity and breadth levels. Overall, the closest the benefits are from the full benefits, the more the cells will tend to green. It is expected that the greater the flexibility, the greater the benefits, so while increasing both CT indicators, rows and columns by policy and instance should tend to a stronger green tone. We see a rupture of these patterns in some cells for some problems, such as instance 40I2 when flexibility is a decision. This happens because some solutions were not optimal in the allowed time with a certain combination of parameters. Later on, we will report on average solution times and gaps.

With the mentioned color scheme, we can observe in Table 5.23 that in all problems when using chaining policies (Best chain, Heuristic, Standard Chain), it was more beneficial to cross-train additional nurses to one additional unit than training them to be able to work in more than one unit. This is perceivable as vertical variations in color are stronger than horizontal variations. For example, in problem 20I1 with the best chain policy, when both CT indicators are 1, it was possible to obtain 67.15% of the total benefits on average. Suppose the number of units where those nurses could work is increased to two. In that case, those benefits raise to 71.62 %, a variation of 5.47%, but instead, if an additional nurse is cross-trained in each unit to work in only one additional unit, those benefits go up to 95.72%, an increase of 28.57%. Different from the chaining policies, when flexibility is a decision, the differences between increasing the intensity or breadth are smaller, and most of the possible benefits (80%) are achieved with fewer levels of intensity and breadth.

Looking for similar insights at Table 5.24 the first thing that catches the eye is that with $H = 0.8$, adding flexibility does not bring any benefit in some problems. This is, no matter the flexibility we add, the cost does not decrease. This was expected as in Table 5.22 for those problems there was no difference in average costs between full and no flexibility policies. Furthermore, the differences are also reduced in all the other problems, showing how the parameter H imposes an important limit to gain benefits from flexibility. In the problems where some benefit was possible, the intensity and breadth behaviors for the chaining policies were similar than with $H = 0.5$, showing greater variations across intensity than breadth, but this time it took more levels of both indicators to bring most of the benefits of flexibility, which suggests that with higher levels of H more investment in flexibility is needed to gain benefits. In some problems with the chaining policies, the analyzed levels of CT indicators were not enough to reach the benefits of full flexibility. In contrast, the policy where flexibility is a decision was able to reach most of the benefits with low levels of cross-training, showing to be more valuable when faced with these restrictive scenarios, where the minimum proportion of time spent in the unit is higher.

We could aggregate the results from the previous tables to compare the standard chain, heuristic, and best chain policies. In a simple overview, among these policies, the differences are relatively small, the best chain presents better results, and the heuristic was not better in every case than the standard chain. To better appreciate the difference among the mentioned chaining policies, we present Table 5.25 that shows the average percentage of flexibility benefits for the best chain, heuristic, and standard chain policies by problem and level of H, as well as the percentage difference between the heuristic and standard chain, and the average benefits by policy for each problem. Differences between the heuristic and standard chain were, on average, slight in each problem, with most of them varying between 0.05% and almost 4%, with exception of problem 30I3 with $H = 0.8$ where the difference was 11.76%. On average, the differences between heuristic and standard chain were 0.72% and 2.35% for each respective H, with the heuristic showing slightly better results in most of the problems.

To have a better look at the number of cross-trainings with the percentage of benefits in each problem and appreciate the differences between policies, we present Figures 5.6 and 5.7, once again for each respective H. In the figures, for each problem, we plotted the average percentage of benefits in each policy versus the number of cross-trainings. These number of cross-trainings can be easily calculated as $Intensity * Breadth * |J|$ for any instance. The purple and green points in each subplot indicate the no flexibility and full flexibility cases, where the minimum and maximum benefits are achieved. The orange, blue, red, and brown lines respectively correspond to flexibility as a decision, best chain, heuristic, and standard chain policies.

For the first level of H in all problems (see Figure 5.6), all the policies managed to obtain most of the benefits (>80%) without needing even a third of the possible cross-trainings. The flexibility as a

Table 5.23: Average flexibility benefits by problem, policy, intensity and breadth , $H = 0.5$

Problem	Breadth/	Best Chain					Flex. As Decision					Heuristic					Standard chain				
	Intensity	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
20I1	1	67.15%	71.62%	71.70%			96.07%	99.64%	99.75%			66.64%	71.47%	71.70%			44.23%	71.18%	71.70%		
20I1	2	95.72%	97.52%	97.52%			99.77%	99.96%	99.92%			94.98%	97.52%	97.52%			91.93%	97.52%	97.52%		
20I1	3	99.92%	99.97%	99.98%			99.81%	99.94%	100.00%			99.80%	99.98%	99.98%			99.33%	99.98%	99.98%		
20I1	4	100.00%	100.00%	100.00%			99.89%	100.00%	100.00%			100.00%	100.00%	100.00%			100.00%	100.00%	100.00%		
20I2	1	64.75%	78.48%	81.54%	81.75%		92.67%	100.00%	100.00%	100.00%		56.71%	75.37%	80.87%	81.75%		58.13%	73.80%	80.18%	81.75%	
20I2	2	89.36%	98.58%	98.64%	98.64%		100.00%	100.00%	100.00%	100.00%		83.07%	98.57%	98.64%	98.64%		88.97%	98.58%	98.64%	98.64%	
20I3	1	57.96%	71.97%	75.48%	76.40%		87.78%	98.88%	99.89%	100.00%		52.28%	68.96%	74.35%	76.40%		42.01%	62.97%	74.27%	76.40%	
20I3	2	82.53%	94.75%	95.58%	95.58%		99.11%	99.95%	100.00%	100.00%		78.18%	94.46%	95.30%	95.58%		79.37%	89.24%	95.19%	95.58%	
20I3	3	93.59%	99.27%	99.29%	99.29%		100.00%	99.98%	99.89%	100.00%		91.07%	99.23%	99.29%	99.29%		93.41%	96.65%	99.29%	99.29%	
30I1	1	37.92%	45.91%	48.44%			73.57%	91.07%	97.22%			36.20%	45.44%	48.44%			30.88%	43.19%	48.44%		
30I1	2	61.47%	71.42%	72.48%			91.05%	99.59%	100.55%			58.95%	70.67%	72.48%			59.86%	70.66%	72.48%		
30I1	3	78.09%	85.91%	86.16%			97.99%	100.69%	100.50%			74.60%	85.99%	86.21%			76.95%	86.01%	86.19%		
30I1	4	89.19%	93.76%	93.80%			100.13%	100.49%	100.26%			85.46%	94.12%	93.90%			88.21%	94.07%	93.90%		
30I1	5	95.26%	97.79%	98.18%			100.26%	100.49%	100.32%			92.99%	98.52%	98.15%			95.58%	98.43%	98.18%		
30I2	1	34.37%	46.69%	51.31%	52.62%		77.70%	98.22%	100.00%	100.00%		32.43%	45.20%	51.01%	52.62%		23.32%	39.94%	50.99%	52.62%	
30I2	2	60.17%	74.69%	76.27%	76.28%		98.53%	100.00%	100.00%	100.00%		57.10%	74.49%	76.03%	76.28%		54.86%	73.71%	76.20%	76.28%	
30I2	3	77.86%	88.63%	88.74%	88.74%		100.00%	100.00%	100.00%	100.00%		73.94%	88.56%	88.74%	88.74%		74.91%	88.62%	88.74%	88.74%	
30I3	1	47.63%	59.44%	61.99%	62.32%	62.33%	81.82%	99.33%	100.00%	100.00%	100.00%	42.98%	56.39%	59.88%	62.14%	62.33%	32.27%	47.72%	60.93%	62.23%	62.33%
30I3	2	75.13%	86.56%	86.89%	86.89%	86.89%	99.35%	100.00%	100.00%	100.00%	100.00%	69.70%	86.26%	86.89%	86.89%	86.89%	64.68%	80.77%	86.49%	86.89%	86.89%
30I3	3	89.27%	95.66%	95.66%	95.66%	95.66%	100.00%	100.00%	100.00%	100.00%	100.00%	84.92%	95.66%	95.66%	95.66%	95.66%	88.78%	95.45%	95.66%	95.66%	95.66%
40I1	1	47.59%	67.70%	75.27%			89.23%	99.94%	100.00%			45.04%	66.55%	75.27%			43.74%	66.05%	75.27%		
40I1	2	76.08%	95.21%	96.63%			99.92%	100.00%	100.00%			73.65%	94.98%	96.63%			75.62%	95.08%	96.63%		
40I1	3	91.28%	99.73%	99.73%			100.00%	100.00%	100.00%			89.67%	99.73%	99.73%			91.26%	99.73%	99.73%		
40I1	4	97.81%	100.00%	100.00%			100.00%	100.00%	100.00%			97.21%	100.00%	100.00%			97.83%	100.00%	100.00%		
40I1	5	99.73%	100.00%	100.00%			100.00%	100.00%	100.00%			99.54%	100.00%	100.00%			99.76%	100.00%	100.00%		
40I1	6	100.00%	100.00%	100.00%			100.00%	100.00%	100.00%			99.97%	100.00%	100.00%			100.00%	100.00%	100.00%		
40I1	7	100.00%	100.00%	100.00%			100.00%	100.00%	100.00%			100.00%	100.00%	100.00%			100.00%	100.00%	100.00%		
40I1	8	100.00%	100.00%	100.00%			100.00%	100.00%	100.00%			100.00%	100.00%	100.00%			100.00%	100.00%	100.00%		
40I2	1	51.22%	69.97%	74.76%	75.79%		89.96%	96.70%	91.85%	98.58%		48.80%	67.74%	73.97%	75.79%		47.88%	68.44%	71.68%	75.79%	
40I2	2	79.13%	93.98%	94.76%	94.80%		91.51%	97.34%	99.59%	91.53%		76.24%	93.85%	94.73%	94.79%		75.29%	92.19%	94.67%	94.79%	
40I2	3	92.22%	99.29%	99.39%	99.44%		99.66%	94.73%	86.55%	94.38%		89.96%	99.39%	99.42%	99.42%		90.70%	98.61%	99.44%	99.42%	
40I2	4	97.87%	100.13%	100.17%	100.32%		93.66%	88.59%	91.68%	83.36%		96.66%	100.46%	100.42%	100.31%		98.16%	100.46%	100.42%	100.33%	
40I2	5	99.71%	100.39%	100.30%	100.35%		100.24%	93.15%	97.79%	98.70%		99.37%	100.48%	100.49%	100.42%		100.27%	100.55%	100.52%	100.41%	
40I2	6	100.21%	100.19%	100.21%	99.66%		100.22%	94.46%	94.26%	96.09%		100.39%	100.41%	100.30%	100.02%		100.64%	100.42%	100.43%	99.69%	
40I3	1	61.88%	73.84%	74.69%	74.79%	74.79%	88.28%	99.93%	100.00%	100.00%	100.00%	54.97%	69.78%	72.70%	74.62%	74.79%	47.10%	64.97%	73.35%	74.74%	74.79%
40I3	2	89.99%	94.39%	94.39%	94.39%	94.39%	99.91%	100.00%	100.00%	100.00%	100.00%	82.60%	94.27%	94.39%	94.39%	94.39%	84.65%	93.87%	94.39%	94.39%	94.39%
40I3	3	98.32%	98.99%	98.99%	98.99%	98.99%	100.00%	100.00%	100.00%	100.00%	100.00%	94.53%	98.99%	98.99%	98.99%	98.99%	98.19%	98.99%	98.99%	98.99%	98.99%
40I3	4	99.91%	99.93%	99.93%	99.93%	99.93%	100.00%	100.00%	100.00%	100.00%	100.00%	98.82%	99.93%	99.93%	99.93%	99.93%	99.91%	99.93%	99.93%	99.93%	99.93%
40I3	5	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.78%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
40I3	6	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.93%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

Table 5.24: Average flexibility benefits by problem, policy, intensity and breadth, $H = 0.8$

Problem	Breadth/	Best Chain					Flex. As Decision					Heuristic					Standard chain				
	Intensity	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
20I1	1	39.95%	43.50%	43.93%			71.06%	93.11%	98.02%			39.54%	43.31%	43.93%			25.24%	42.88%	43.93%		
20I1	2	68.75%	71.92%	71.94%			93.74%	99.21%	99.85%			68.00%	71.87%	71.94%			59.14%	71.86%	71.94%		
20I1	3	86.33%	89.72%	89.75%			97.99%	99.73%	99.94%			85.52%	89.75%	89.73%			74.30%	89.75%	89.73%		
20I1	4	95.30%	97.44%	97.50%			99.28%	99.82%	99.87%			95.01%	97.50%	97.51%			85.33%	97.51%	97.50%		
20I2	1	31.58%	35.56%	36.62%	36.81%		49.89%	82.39%	98.73%	100.00%		25.21%	33.01%	35.29%	36.81%		25.51%	31.96%	35.24%	36.81%	
20I2	2	50.33%	61.17%	64.22%	64.84%		82.39%	100.00%	100.00%	100.00%		42.91%	60.63%	62.44%	64.84%		46.00%	59.54%	61.62%	64.84%	
20I3	1	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%	0.00%	
20I3	2	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%	0.00%	
20I3	3	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%	0.00%	
30I1	1	0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%		
30I1	2	0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%		
30I1	3	0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%		
30I1	4	0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%		
30I1	5	0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%			0.00%	0.00%	0.00%		
30I2	1	16.01%	21.17%	22.99%	23.82%		35.70%	66.46%	86.82%	97.24%		13.92%	19.54%	22.43%	23.82%		10.10%	16.48%	22.41%	23.82%	
30I2	2	28.79%	38.49%	42.66%	44.82%		66.47%	97.27%	100.00%	100.00%		25.76%	37.99%	41.30%	44.82%		23.87%	36.54%	42.35%	44.82%	
30I2	3	40.00%	53.52%	59.73%	62.20%		86.84%	100.00%	100.00%	100.00%		36.58%	52.84%	59.52%	62.20%		35.60%	52.62%	59.30%	62.20%	
30I3	1	24.75%	29.53%	30.50%	30.61%	30.62%	46.20%	76.00%	94.22%	99.97%	100.00%	20.57%	26.84%	28.02%	30.12%	30.62%	16.13%	22.30%	28.87%	30.44%	30.62%
30I3	2	42.36%	51.43%	53.48%	53.72%	53.78%	76.11%	99.97%	100.00%	100.00%	100.00%	36.48%	49.68%	52.75%	53.07%	53.78%	32.93%	44.62%	50.44%	53.52%	53.78%
30I3	3	56.31%	68.70%	71.40%	71.74%	71.76%	94.19%	100.00%	100.00%	100.00%	100.00%	49.88%	67.82%	70.85%	71.70%	71.76%	52.69%	63.91%	69.30%	71.75%	71.76%
40I1	1	18.87%	26.54%	30.81%			40.12%	68.09%	85.95%			16.77%	25.49%	30.81%			16.97%	24.94%	30.81%		
40I1	2	32.46%	46.24%	53.46%			68.17%	96.13%	100.00%			29.94%	45.65%	53.46%			31.19%	45.91%	53.46%		
40I1	3	43.75%	62.00%	70.04%			85.91%	100.00%	100.00%			40.97%	62.21%	70.04%			43.04%	62.21%	70.04%		
40I1	4	53.44%	74.36%	82.07%			96.10%	100.00%	100.00%			50.56%	74.38%	82.07%			52.90%	74.49%	82.07%		
40I1	5	62.01%	84.00%	90.49%			99.51%	100.00%	100.00%			58.67%	84.16%	90.49%			61.35%	84.27%	90.49%		
40I1	6	69.04%	90.91%	95.35%			100.00%	100.00%	100.00%			65.40%	91.18%	95.35%			69.26%	91.18%	95.35%		
40I1	7	74.55%	95.24%	98.40%			100.00%	100.00%	100.00%			70.73%	95.96%	98.40%			75.20%	95.98%	98.40%		
40I1	8	78.71%	97.44%	99.60%			100.00%	100.00%	100.00%			74.65%	98.50%	99.60%			80.10%	98.44%	99.60%		
40I2	1	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%	0.00%	
40I2	2	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%	0.00%	
40I2	3	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%	0.00%	
40I2	4	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%	0.00%	
40I2	5	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%	0.00%	
40I2	6	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%	0.00%	
40I3	1	31.33%	37.04%	37.76%	37.88%	37.92%	45.52%	76.85%	95.06%	99.73%	100.00%	24.82%	33.03%	34.56%	36.91%	37.92%	23.64%	30.66%	35.17%	36.83%	37.92%
40I3	2	51.06%	61.14%	62.51%	62.89%	62.93%	76.98%	99.65%	100.00%	100.00%	100.00%	42.33%	59.28%	61.80%	61.95%	62.93%	44.74%	56.78%	61.20%	62.49%	62.93%
40I3	3	65.27%	77.27%	78.79%	79.04%	79.07%	95.30%	100.00%	100.00%	100.00%	100.00%	55.94%	76.36%	78.62%	79.05%	79.07%	63.35%	75.03%	78.43%	78.98%	79.07%
40I3	4	76.26%	87.98%	88.91%	88.97%	88.98%	99.69%	100.00%	100.00%	100.00%	100.00%	66.53%	87.91%	88.85%	88.98%	88.98%	77.07%	87.06%	88.69%	88.98%	88.98%
40I3	5	83.74%	94.17%	94.63%	94.64%	94.64%	100.00%	100.00%	100.00%	100.00%	100.00%	74.12%	94.29%	94.64%	94.64%	94.64%	86.27%	93.84%	94.63%	94.64%	94.64%
40I3	6	88.03%	97.46%	97.76%	97.76%	97.76%	100.00%	100.00%	100.00%	100.00%	100.00%	79.24%	97.63%	97.76%	97.76%	97.76%	91.50%	97.39%	97.73%	97.76%	97.76%

Table 5.25: Comparison of average flexibility benefits among chaining policies

Problem	H = 0.5				H = 0.8			
	Best Chain	**Heuristic**	**Standard chain**	**Difference**	**Best Chain**	**Heuristic**	**Standard chain**	**Difference**
20I1	91.76%	**91.63%**	89.45%	**2.18%**	74.67%	**74.47%**	70.76%	**3.71%**
20I2	86.47%	84.20%	**84.84%**	-0.64%	47.64%	45.14%	45.19%	-0.05%
20I3	86.81%	**85.37%**	83.64%	**1.73%**	-	-	-	-
30I1	77.05%	76.14%	**76.20%**	-0.06%	-	-	-	-
30I2	68.03%	**67.10%**	65.74%	**1.36%**	37.85%	**36.73%**	35.84%	**0.89%**
30I3	79.20%	**77.86%**	76.16%	**1.70%**	49.38%	**47.60%**	35.84%	**11.76%**
40I1	93.61%	93.25%	**93.36%**	-0.11%	67.91%	66.89%	67.92%	-1.03%
40I2	92.68%	**92.24%**	92.13%	**0.11%**	-	-	-	-
40I3	94.05%	**93.19%**	93.01%	**0.18%**	74.45%	72.28%	73.47%	-1.19%
Average	85.52%	**84.55%**	83.84%	0.72%	58.65%	**57.19%**	54.84%	2.35%

decision policy achieved almost all benefits of full flexibility with at most the second-lowest number of cross-trainings in all problems, whereas the remaining policies did so with a higher number of cross-trainings. For problem 40I2, the flexibility as decision policy struggled to reach all the benefits after 25 cross-trainings, results that can be explained by the presence of important gaps in certain instances of the problem. There are small percentage differences among the chaining policies (less than 10% for most of the required cross-trainings values) in all problems, especially before a certain number of cross-trainings e.g. 8 cross-trainings in problem 20I1, 15 for problem 20I2, and so on, after which the three chaining policies overlap entirely, showing no differences in benefits among them. This overlapping suggests that there is an investment level after which using one or another of the selected chaining policies does not make any difference by bringing additional benefits. However, in some cases and especially with a low number of cross-trainings (before the point when the policies overlap) using the best chain and the heuristic can grasp some additional benefits.

In Figure 5.7 the flexibility as a decision policy results were similar than with an H of 0.5, but in every case, more cross-trainings were necessary to reach most of the flexibility benefits. The chaining policies could not reach all of these benefits with the given cross-trainings in problems 20I2, 30I2, and 30I3, and in problems 20I1, 40I1, and 40I2 the full benefits were reached with more than 60% of the total cross-trainings. The differences among chaining policies were non-existent from the number of cross-trainings where the policies overlapped, and were at most 20% with lower number of cross-trainings. See for example, for problem 20I1 from 24 cross-trainings the average benefits are the same for the different chaining policies, and before that number the heuristic and best chain show some additional benefits over the standard chain. The best chain kept showing an improvement over the heuristic and standard chain in all problems in low cross-training levels, but the differences between these last two policies were rather small (less than 5%) in most of cross-trainings levels and problems, and favored the heuristic in some cases and the standard in others.

We can also analyze the influence of the H levels in cost. We noted before that with a higher level of H more cross-trainings are needed to reach most of the flexibility benefits, and those benefits are reduced in some problems with respect to the 0.5 level. In Figure 5.8 we can see the average cost of all combinations of intensity, breadth, and demand for each problem classified by H. The percentages on top of the $H = 0.8$ bars indicate the percentage increment from the average cost when $H = 0.5$. As it can be observed, the increments varied a lot across problems, ranging from 4% in problem 20I2 to more than 400% in instance 40I2.

As mentioned before, not all instances in all problems were optimally solved. This can be verified

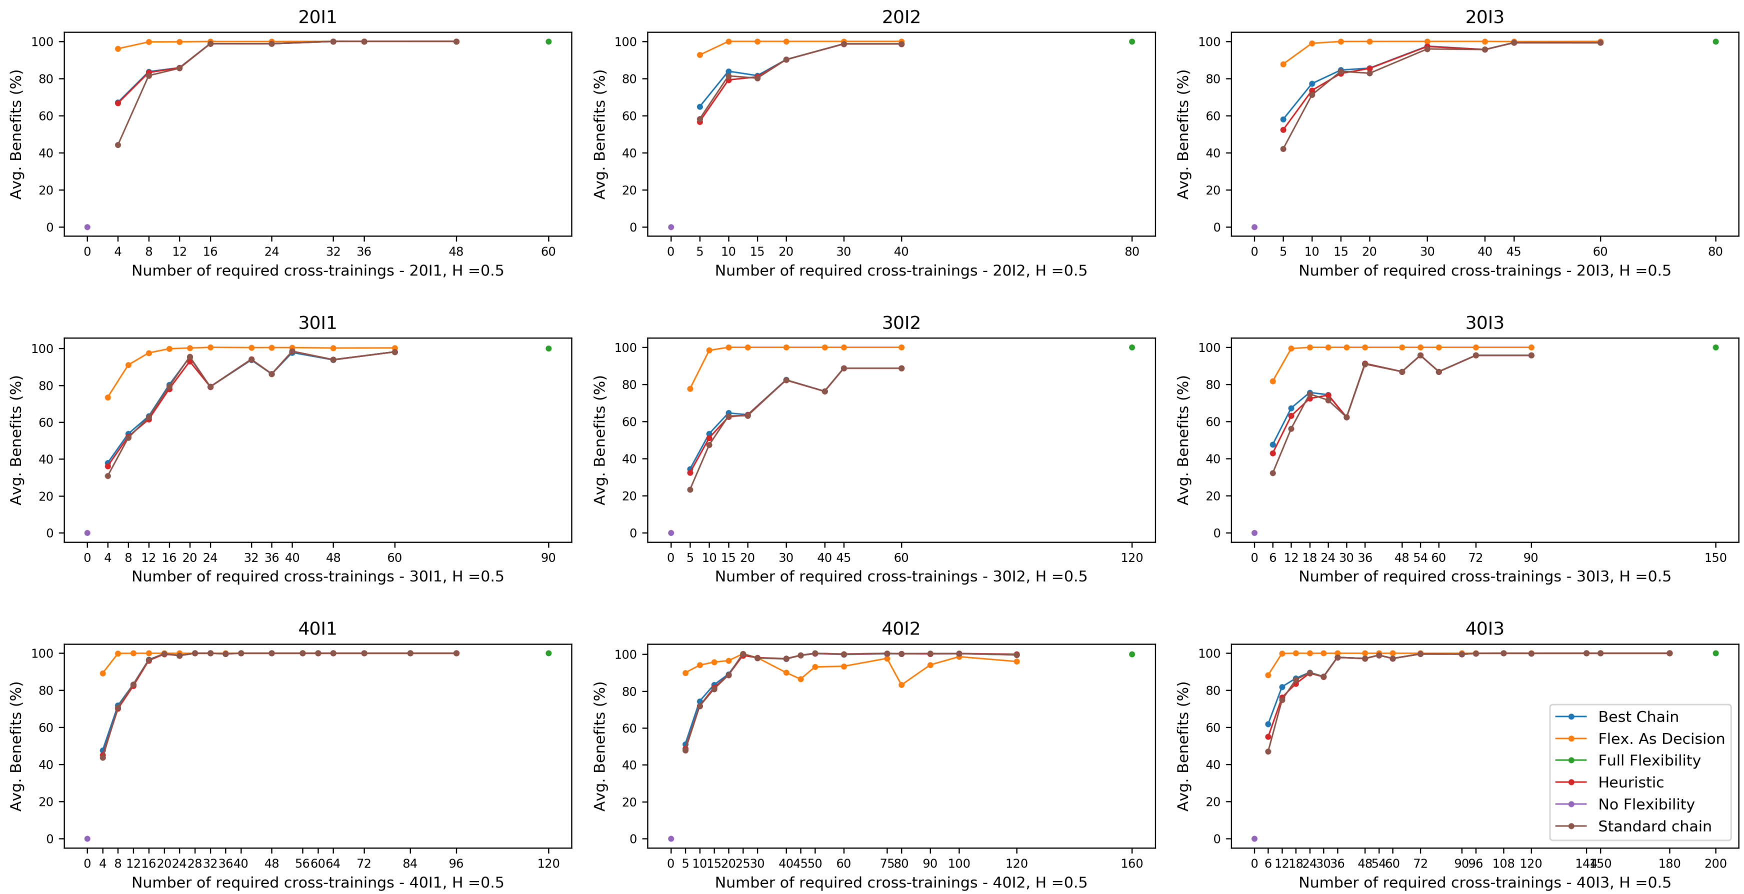

Figure 5.6: Avg. Benefits (%) vs Number of required cross-trainings by problem, $H = 0.5$

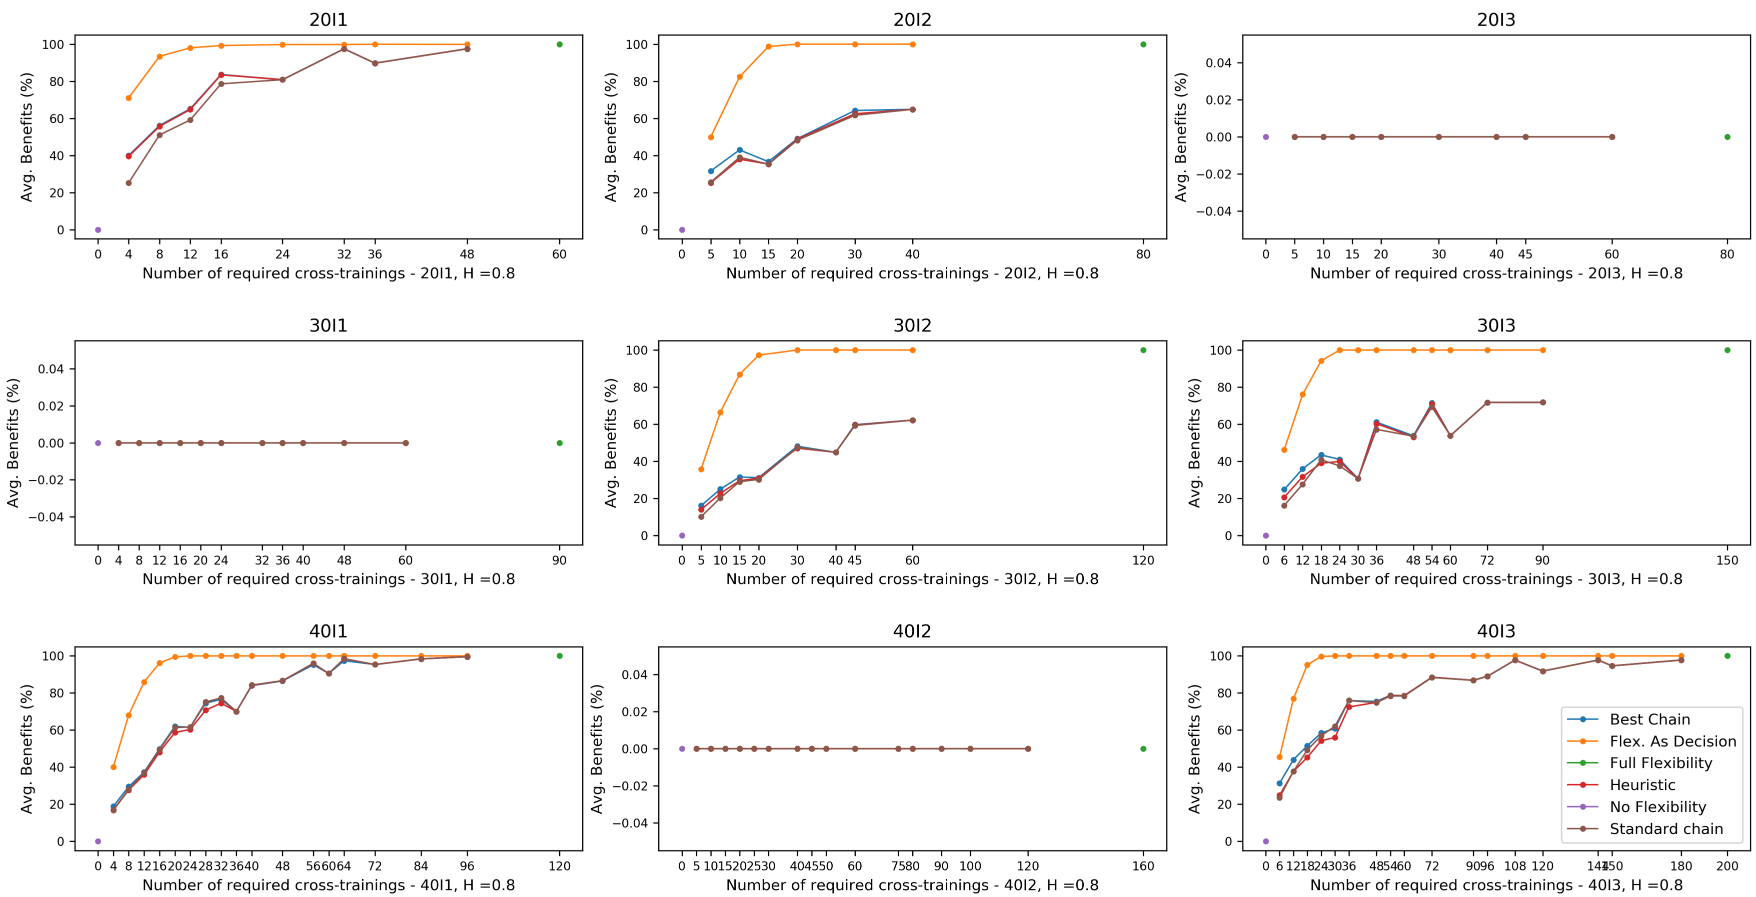

Figure 5.7: Avg. Benefits (%) vs Number of required cross-trainings by problem, $H = 0.8$

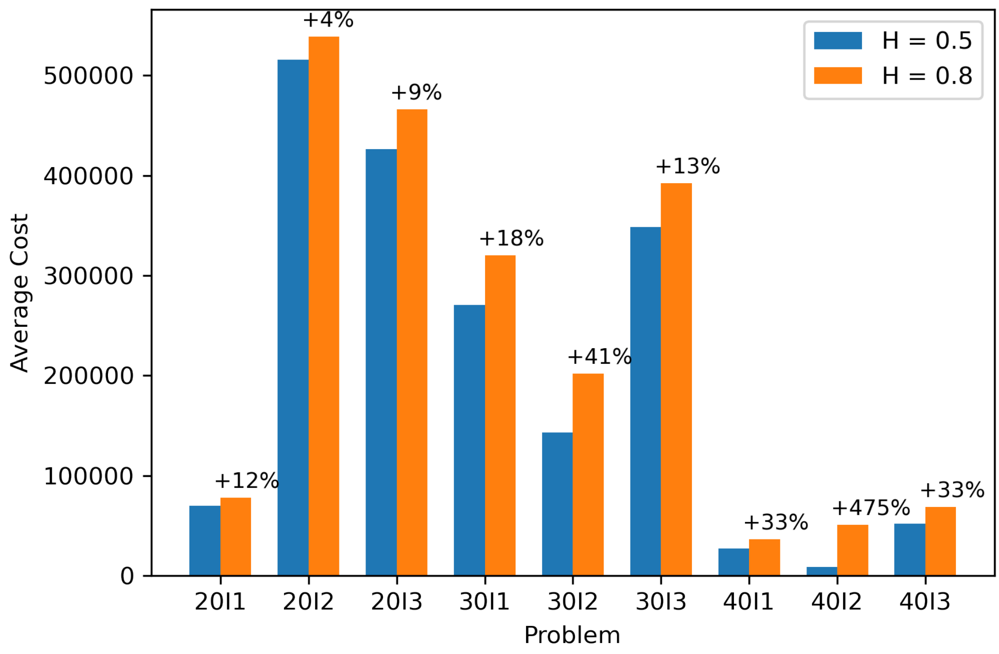

Figure 5.8: Average cost by problem and H

by the presence of gaps in some instances, the most remarkable one being problem 40I2, where in the results table we see some alterations in the color patterns. Taking a look at the box-plot in Figure 5.9 we first note that the median (orange line) of the average gaps for every combination of problem and H is close to zero. In almost all problems, there were outliers, where higher gaps are present. In most of the cases, they were lower than 5%, but some high gaps are observed, the most representatives being in problem 40I2 with $H = 0.5$ where there were gaps over 20%.

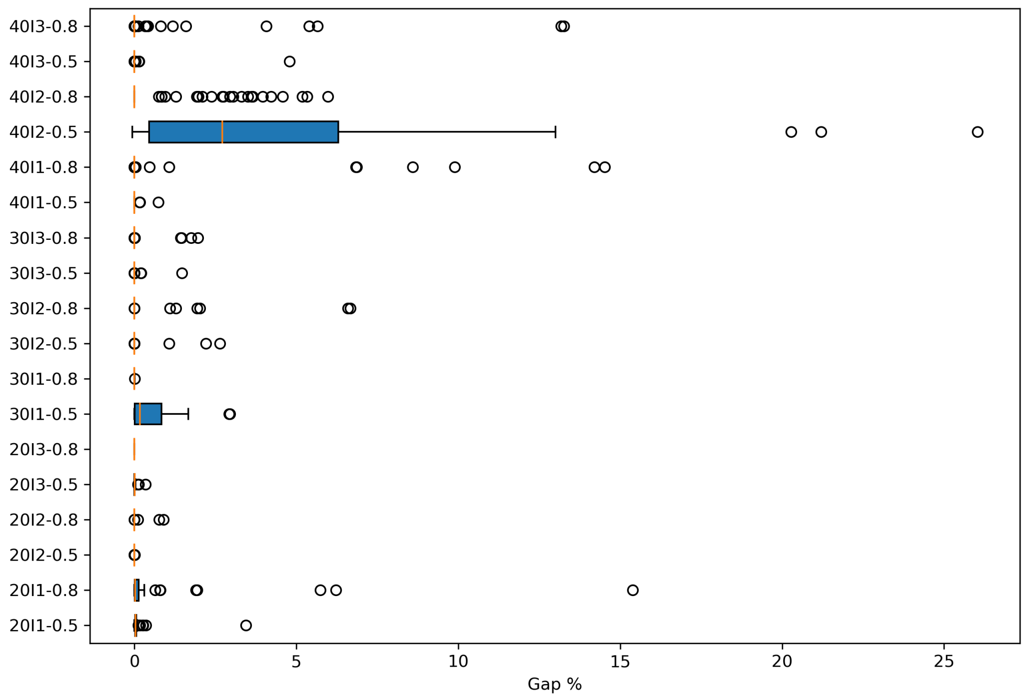

Figure 5.9: Average gaps by problem and H

We could also distinguish these gaps by policy for each level of H, as respectively seen in Figures 5.10 and 5.11. In these figures, we note how the flexibility as a decision had a wider range of gaps when compared with the other policies in both H levels. It is also noticeable that with $H = 0.8$ the outliers are less variable in most of the policies. The previous gap analysis could explain the jumps in Figure 5.6 for instance 40I2, where the flexibility as a decision policy reaches all the benefits, and then with more flexibility struggles to maintain those benefits.

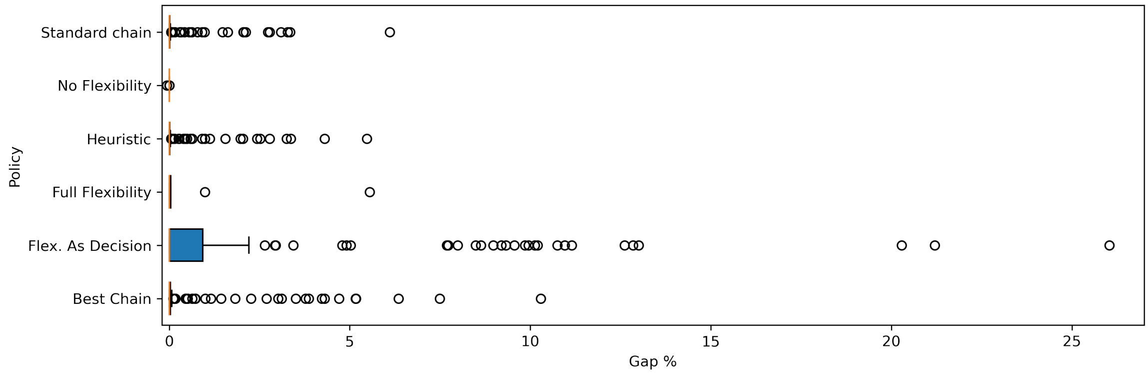

Figure 5.10: Gaps percentage by policy, $H = 0.5$

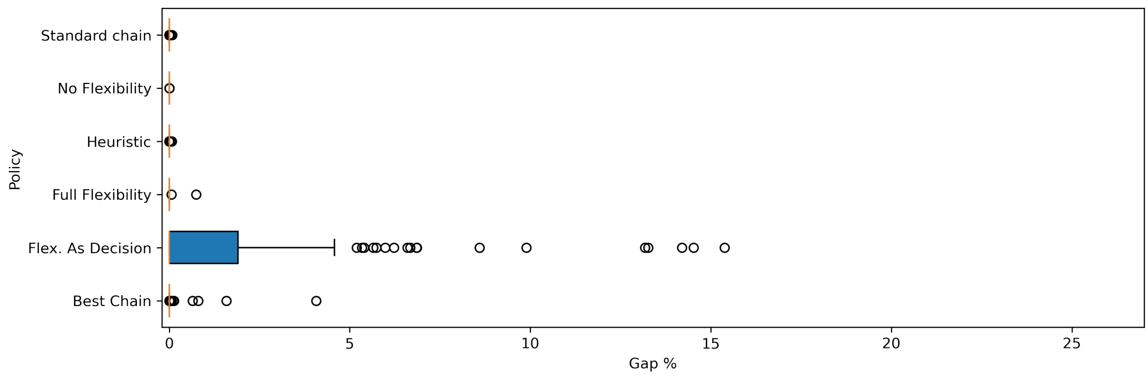

Figure 5.11: Gaps percentage by policy, $H = 0.8$

When looking similarly at times by policy for each H level in 5.12 and 5.13, we see wider ranges in the flexibility as a decision policy, the only policy that had instances of the problem run using most of the allowed time, but in close to 75% of the cases it used less than around 400s for $H = 0.5$ and 300s for $H = 0.8$. There is also a difference in the means of both H levels, suggesting lower solution times in half of the instances solving times when $H = 0.5$. When compared, there was also a wider range of times in the standard chain, heuristic, and best chain with the least restrictive proportion of time in the home unit.

5.3 Experiments for the investment in flexibility

The previous section's experiments gave a general idea of the existing dynamic between the benefits of flexibility and investment. When plotting the average benefits vs. the number of cross-trainings, we saw that it is not necessary to fully invest in flexibility (cross-train all workers to all units) to obtain all its benefits. It was also observed that the flexibility as a decision policy

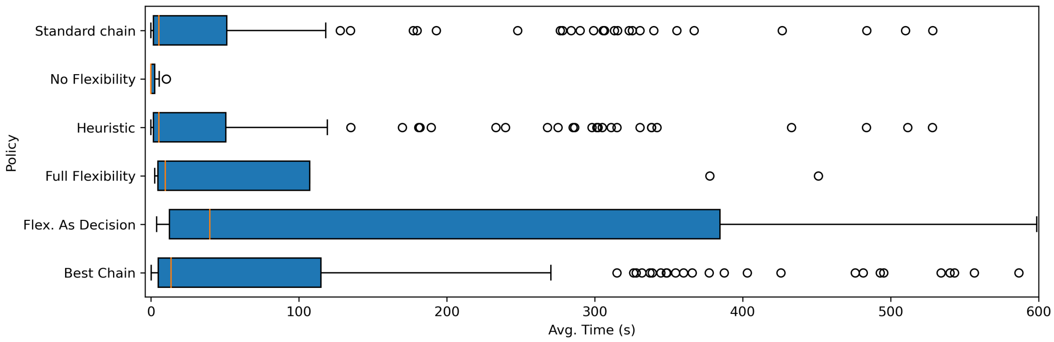

Figure 5.12: Average times by policy, $H = 0.5$

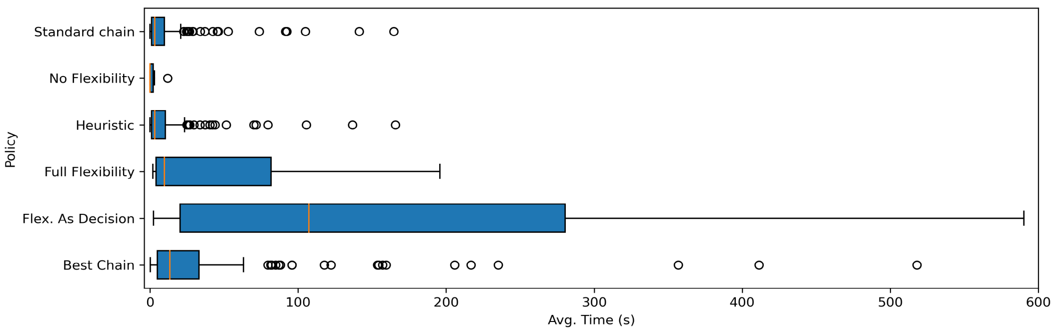

Figure 5.13: Average times by policy, $H = 0.8$

made better use of the cross-trainings, achieving a higher percentage of benefits when compared with other policies at different investment limits. In the experiments of this section, we will fix the number of cross-trainings in a consecutive way such that it is easier to appreciate the trade-off of total costs and the investment in flexibility. For this experiment, we selected instances 20I1, 30I1, and 40I1, an H of 0.5, and the set of 100 demands from the previous experiment. Only one problem per size was selected as it is expected that the trade-off will be similarly perceivable in all problems. The lowest H was selected because as the previous experiment showed, more benefits are possible as this parameter is lower. This time, we will change the inequality in constraint (4.18) to an equality, and the DC will take values (investment) from 0 to 24 to see how the cost changes by increments of one cross-training. Once again, each instance of the problems will run for a limit of 10 minutes.

Results discussions for investment in flexibility

In Figures 5.14 to 5.16, we see results of this experiment for problems 20I1, 30I1, and 40I1, respectively. In the vertical axis of each figure, we see the total cost. The blue, orange, and green points and lines correspond to the maximum, average, and minimum total cost of the 100 demands when using the specific level of investment. The horizontal axis shows the percentage of total cross-trainings used, corresponding to the percentage that the consecutive investments represent for each problem. For example, in problem 20I1, the maximum number of cross-trainings (maximum investment) is 60, which corresponds to cross-train each nurse to every other unit, and one cross-training represents 1.7% (1/60) of that maximum possible investment; consequently, 2 cross-trainings are 3.3% and so on.

The case where the percentage of investment is zero in the figures serves as a reference because it indicates the maximum, average, and minimum cost when there is no investment in flexibility. For example, in problem 20I1 (see Fig. 5.14), with no flexibility, in the worst-case scenario among the evaluated demands, the maximum cost can reach almost 200000 and the best scenario (minimum cost) around 58000. In the three plots, it is notorious how with the increase of the investment, the average total costs reduce, but these reductions diminish over the increments of the investment until a level where the total cost seems to converge to a value. For problem 20I1, around 10% of the total possible cross-trainings are needed to achieve most of the cost reduction, in problem 30I1 this value is around 16.7% , and for problem 40I1 around 5.8%. Those percentages correspond to 6, 16, and 7 cross-trainings for each respective problem.

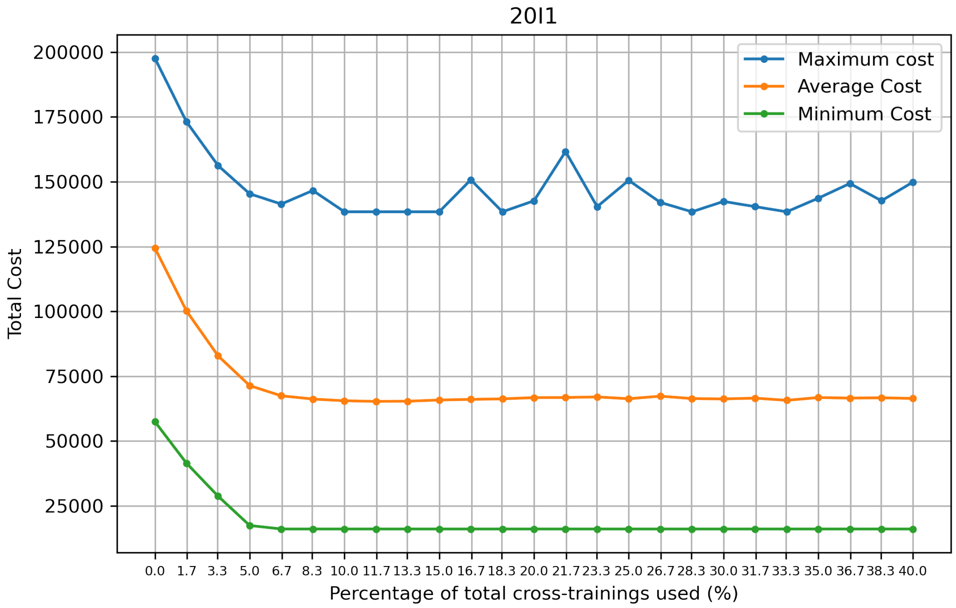

Figure 5.14: Total Cost vs. Percentage of total cross-trainings used for problem 20I1

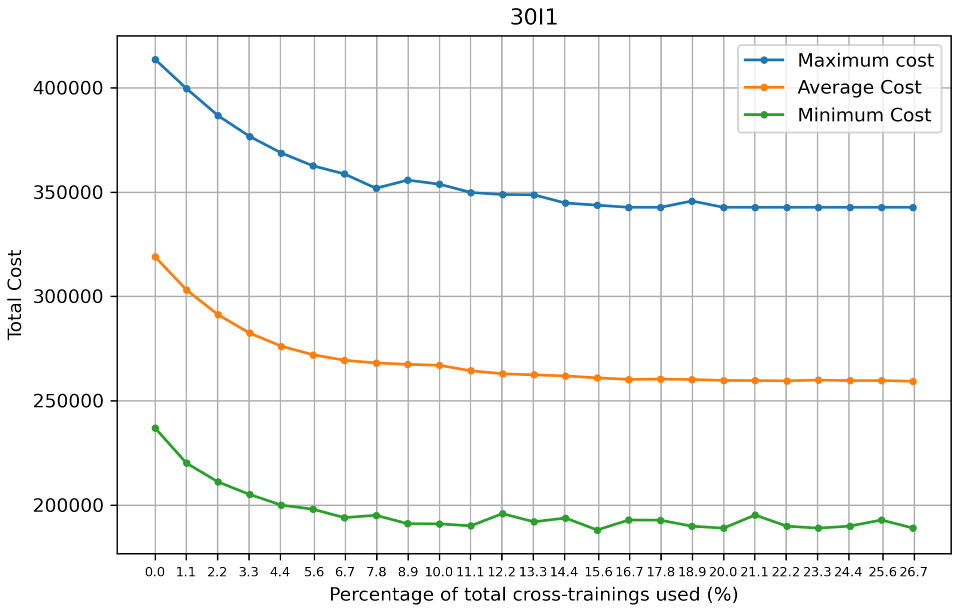

Figure 5.15: Total Cost vs. Percentage of total cross-trainings used for problem 30I1

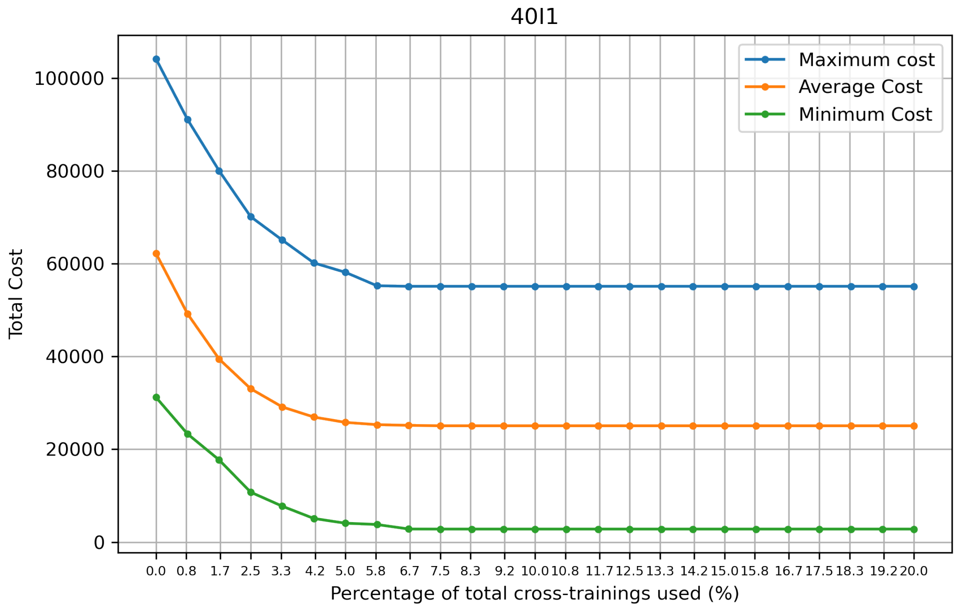

Figure 5.16: Total Cost vs. Percentage of total cross-trainings used for problem 40I1

We could also analyze the average proportion of the time worked by dedicated nurses in every unit. This might not have high importance in the current flexibility model as cross-trained nurses are considered to be equally efficient than dedicated nurses when floating to another units, trainings are considered 100% effective and learning with time is not considered. However, if this is not the case and dedicated nurses are considered to provide a greater quality of care than cross-trained ones it might be more beneficial to have dedicated nurses working most of the time in every unit. Note that by itself, the continuity of care parameter (H) does not guarantees this happens, as it only considers regular working time (not overtime), and it limits the individual time of nurses in other units and not the total time that all cross-trained nurses can work in each unit.

In Figure 5.17 we show the averaged proportion of dedicated working time versus the number of cross-trainings for each unit in problem 20I1, as well as the average total costs. In the Figure, we see how the total cost reduces with the increase of flexibility, as seen before, and the proportion of the total working time in each unit by its dedicated nurses also reduces. For example, in problem 20I1 with the investment of 24 cross-trainings, this average proportion reduces to almost 0.5 in units 3 and 4, meaning that half of the time, patients are attended by cross-trained nurses in those units. This proportion seems to reduce even more as flexibility increases, which is explained by the fact that cross-trained workers might be interchangeably used for dedicated ones, and more of those cross-trained workers are allocated in the schedule of float units. With the minimal approximate number of cross-trainings that minimizes the average total cost for problem 20I1 (6 in the Figure), the proportion of dedicated working time reduces to around 0.8 for units 3 and 4 and maintains close to 1 in units 1 and 2, which suggests that the latter pair of units float nurses to the former pair. This is consistent with the problem as the first two units have more dedicated nurses, six each, than units 3 and 4.

Similarly, in Figure 5.18 for problem 30I1 with 16 cross-trainings, three of the four units (units 1,2, and 4) would use dedicated workers in their schedule at least 80% of the time, and unit 3 almost 70% of the time. In Figure 5.19 for problem 40I2, the reduction seems to be less, as with

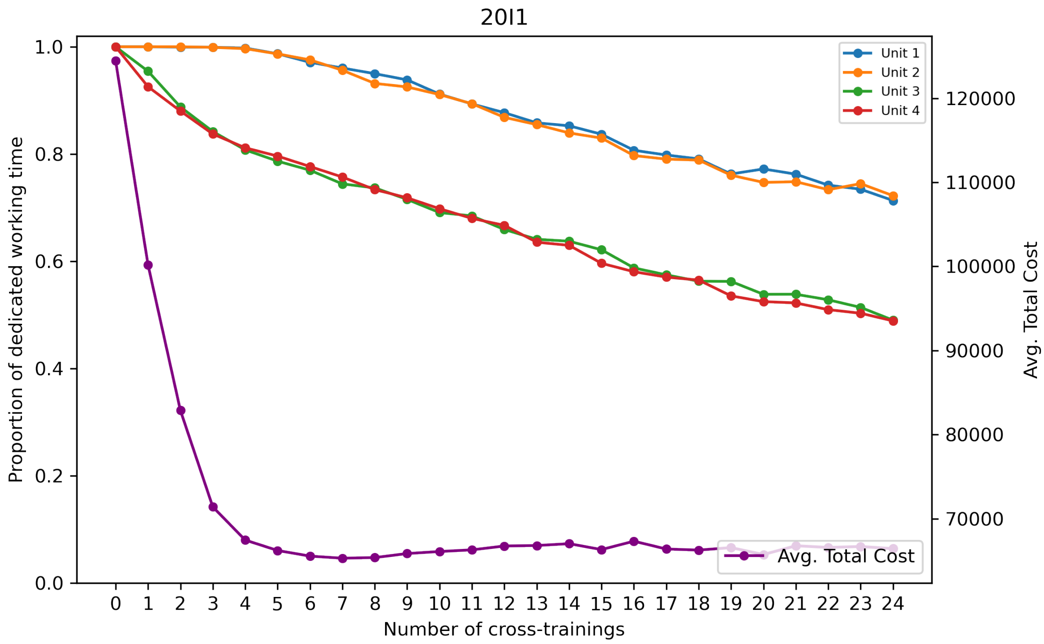

Figure 5.17: Proportion of dedicated working time and Avg. total cost vs. number of cross-trainings - 20I1

7 cross-trainings the proportion of time for units from 1 to 3 is close to 0.9 and a almost 0.8 for unit 4. If we look at the reduction of the proportion of dedicated working time with the increase of flexibility in relation to the distribution of nurses in all three problems, we note that the greater is the number of nurses in a unit the proportion of time is also grater for any investment. We can also see that having higher flexibility than required can be negative in terms of quality of service for some units if nurses are not considered equally efficient in float units.

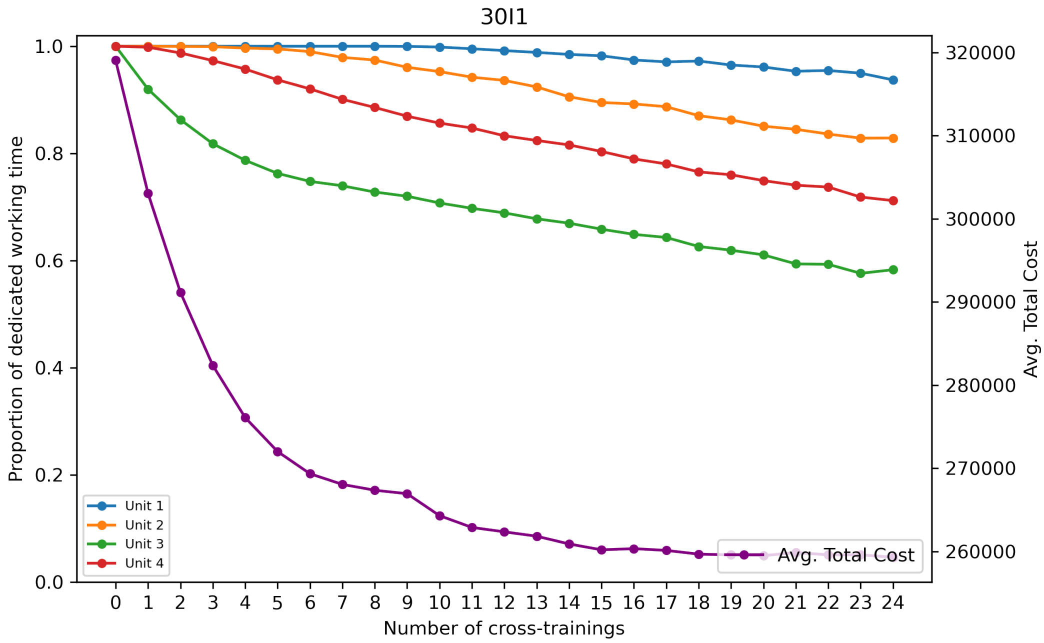

Figure 5.18: Proportion of dedicated working time and Avg. total cost vs. number of cross-trainings - 30I1

Once again, gaps were present in the instances of every problem. Contrary to the intuition

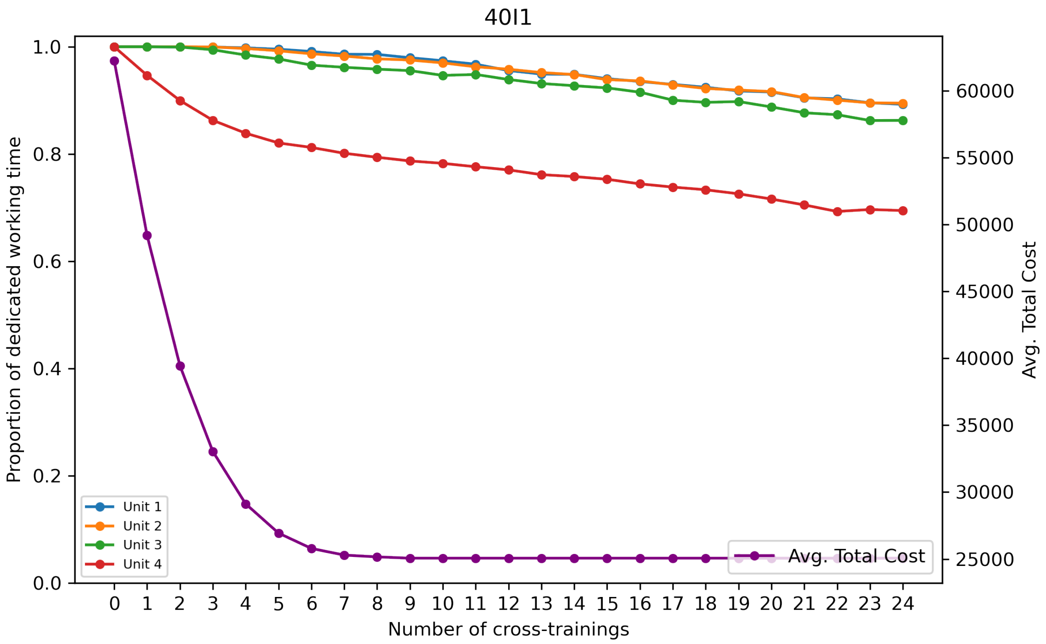

Figure 5.19: Proportion of dedicated working time and Avg. total cost vs. number of cross-trainings - 40I1

that an increased number of nurses would imply higher gaps, despite being the problem with fewer nurses, problem 20I1 was the one that showed to have higher average gaps in most of the different number of cross-trainings investment levels. It stands out an average gap of almost 12% when investing two cross-trainings in the aforementioned problem, and gaps up to 2% for the remaining problems across all number of cross-trainings. Regarding average times, that can be seen in Figure 5.21 problems 20I1 and 30I1 used on average almost all the available time to solve, while problem 40I1 required less than 400s in most of the number of cross-trainings, being the highest average time a little over 500s with five cross-trainings.

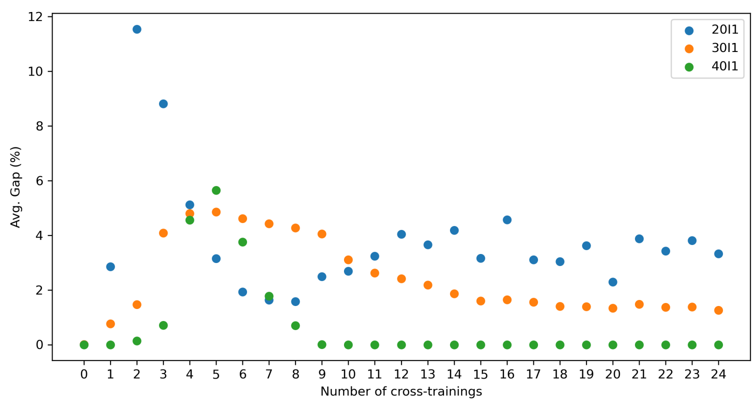

Figure 5.20: Average gaps by number of cross-trainings and problem

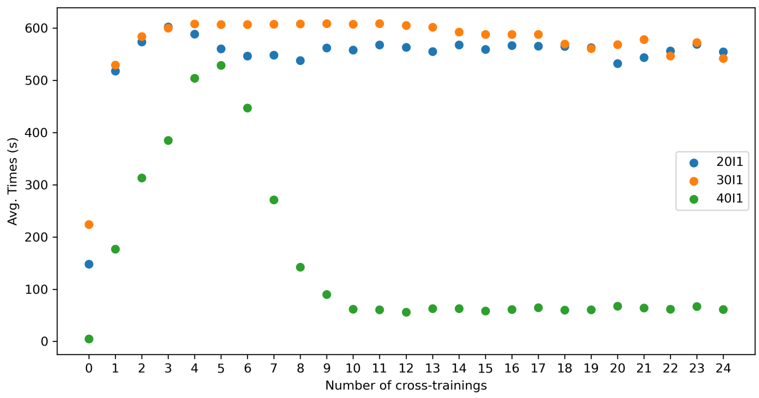

Figure 5.21: Average times by number of cross-trainings and problem

CHAPTER 6

Final remarks and future research directions

In this work, we proposed variations for a mid-term nurse scheduling problem to analyze the benefits of flexibility. The first variation considers the possibility to invest in flexibility, and the second finds the best chain. In addition, we developed a heuristic that designs a chain based on the demands of every unit. Preliminary experiments were performed with different intensities, minimum proportions of time in the home unit (H), demands, and policies. We observed that it is not necessary to fully invest in flexibility to obtain most of its benefits and that the flexibility as a decision and chaining policies (Best chain, heuristic, standard chain) showed better results in terms of savings.

From the preliminary experiments, we selected the best-performing policies and tested them in nine problems to see if their performance was maintained when faced with different levels of intensity, breadth, H, and demands. When analyzing intensity and breadth, the results showed that investing in intensity can bring more benefits than investing in breadth; training two nurses each one to an additional unit is more beneficial than training only one of them to two units. In other words, the intensity has more substantial benefits than breadth, an insight that was found in the works of Fügener et al. (2018) and Wright and Bretthauer (2010), and again observed in our results for the proposed problems where it was possible to benefit from cross-training flexibility. Furthermore, the minimum proportion of working time of every nurse in its home unit (H) had a strong impact on the average cost of every problem, and therefore in the benefits of flexibility, so much that with $H = 0.8$ it was not possible to obtain average benefits in three problems. The possible average benefits were reduced in the remaining problems, and it took more levels of breadth and intensity to reach most or all benefits.

After making a simple comparison on the average percentage benefits of the chaining policies (best chain, standard chain, and heuristic), first, we found that, as expected, the proposed best chain shows greater average benefits than the other two chaining policies. Then, we observed that the heuristic could grasp some additional average benefits in most problems when $H = 0.5$. However, this did not hold when $H = 0.8$ because, from the six problems where it was possible to benefit from flexibility, half showed slightly better results for the heuristic and half for the standard chain. Plotting the results for every H and problem versus the number of cross-trainings for the selected policies showed that differences among the chaining policies, if existent, were small and occurred in most problems with a low number of cross-trainings after which considering the order of units with the best chain or heuristic would not bring any additional benefits. This shows the robustness of the chaining policy, even when faced with heterogeneous systems, like in this case, where the capacities of the units differed.

In experiment two, the flexibility as a decision policy was capable of achieving most of the benefits of flexibility with fewer cross-training than the chaining policies in every problem. It showed to be even more valuable when $H = 0.8$ where the chaining policies took more number of cross-trainings to reach most of the benefits. A clear drawback from this policy seems to be its complexity and size, augmented by the new constraints and variables, which translated in average higher gaps ranges and times. In experiment 3, with different levels of investment in the three selected problems, less than 20% of the total possible investment in flexibility was necessary to achieve all its benefits. On the other hand, we took a look at the proportion of dedicated working time by unit, a measure that we could relate to the quality of service, in the case when cross-training is not entirely effective or cross-trained nurses are not completely efficient in float units. We observed the trade-off between investment and this proportion of time and noted how an increase of the former caused a reduction in the latter for the different units in all problems. The previous trade-off could be of value for decision-makers who wish to know the investment level of flexibility at which the benefits are higher and limit the proportion of total working time that cross-trained nurses work in every unit. The results also suggest that having more flexibility than necessary, in addition, to not add any benefit in terms of cost could have a negative effect as the time of dedicated time in units reduces, which could led to a reduction in on the quality of service, when depth is less that 100%.

Future works could include a new parameter that represents how cross-trained nurses are less efficient when working in float units and analyze the influence of this parameter in the final schedule decision. Also, as shown in the results, the flexibility as a decision model presented high gaps and times in some instances of the problems, which encourages to find: a) alternative ways of modeling the problem, or include additional cuts, given the awareness that the proposed approach might have problems with symmetry issues, and b) explore new solution methods to solve large problems, where more nurses, units, and shifts can be considered. Finally, some remarks from the base model paper can be included, like considering absenteeism and better ways of modeling demands like by skill classes.

Bibliography

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

S. R. Agnihothri and A. K. Mishra. Cross-training decisions in field services with three job types and server–job mismatch. *Decision Sciences*, 35(2):239–257, 2004.

S. Andradóttir, H. Ayhan, and D. G. Down. Design principles for flexible systems. *Production and Operations Management*, 22(5):1144–1156, 2013.

J. F. Bard and H. W. Purnomo. Short-term nurse scheduling in response to daily fluctuations in supply and demand. *Health Care Management Science*, 8(4):315–324, 2005.

M. J. Brusco. An exact algorithm for a workforce allocation problem with application to an analysis of cross-training policies. *Iie Transactions*, 40(5):495–508, 2008.

M. J. Brusco and T. R. Johns. Staffing a multiskilled workforce with varying levels of productivity: An analysis of cross-training policies. *Decision sciences*, 29(2):499–515, 1998.

M. J. Brusco, T. R. Johns, and J. H. Reed. Cross-utilization of a two-skilled workforce. *International Journal of Operations & Production Management*, 18(6):555–564, 1998.

B. Büke, Ö. M. Araz, and J. W. Fowler. Cross-training with imperfect training schemes. *Production and Operations Management*, 25(7):1216–1231, 2016.

E. K. Burke, P. De Causmaecker, G. V. Berghe, and H. Van Landeghem. The state of the art of nurse rostering. *Journal of scheduling*, 7(6):441–499, 2004.

G. M. Campbell. A two-stage stochastic program for scheduling and allocating cross-trained workers. *Journal of the Operational Research Society*, 62(6):1038–1047, 2011.

G. Chauhan and T. Singh. Lean manufacturing through management of labor and machine flexibility: a comprehensive review. *Global Journal of Flexible Systems Management*, 12(1-2):59–80, 2011.

G. Chauhan and T. Singh. Development and validation of resource flexibility measures for manufacturing industry. *Journal of Industrial Engineering and Management*, 7(1):21–41, 2014.

P. De Causmaecker and G. V. Berghe. A categorisation of nurse rostering problems. *Journal of Scheduling*, 14(1):3–16, 2011.

D. R. De Oliveira, R. H. Griep, L. F. Portela, and L. Rotenberg. Intention to leave profession, psychosocial environment and self-rated health among registered nurses from large hospitals in Brazil: A cross-sectional study. *BMC Health Services Research*, 17(1):1–10, 2017. ISSN 14726963. doi: 10.1186/s12913-016-1949-6. URL http://dx.doi.org/10.1186/s12913-016-1949-6.

F. F. Easton. Cross-training performance in flexible labor scheduling environments. *Iie Transactions*, 43(8):589–603, 2011.

A. T. Ernst, H. Jiang, M. Krishnamoorthy, and D. Sier. Staff scheduling and rostering: A review of applications, methods and models. *European journal of operational research*, 153(1):3–27, 2004.

D. J. Fiorotto, R. Jans, and S. A. de Araujo. Process flexibility and the chaining principle in lot sizing problems. *International Journal of Production Economics*, 204:244–263, 2018.

A. Fügener, A. Pahr, and J. O. Brunner. Mid-term nurse rostering considering cross-training effects. *International Journal of Production Economics*, 196:176–187, 2018.

A. Gnanlet and W. G. Gilland. Sequential and simultaneous decision making for optimizing health care resource flexibilities. *Decision Sciences*, 40(2):295–326, 2009.

A. Gnanlet and W. G. Gilland. Impact of productivity on cross-training configurations and optimal staffing decisions in hospitals. *European Journal of Operational Research*, 238(1):254–269, 2014.

S. D. Hartog. On the complexity of nurse scheduling problems. *Master's Thesis*, 4(2):71–74, 2016.

C. A. Henao, J. C. Muñoz, and J. C. Ferrer. The impact of multi-skilling on personnel scheduling in the service sector: a retail industry case. *Journal of the Operational Research Society*, 66(12): 1949–1959, 2015.

C. A. Henao, J. C. Ferrer, J. C. Muñoz, and J. Vera. Multiskilling with closed chains in a service industry: A robust optimization approach. *International Journal of Production Economics*, 179: 166–178, 2016.

C. A. Henao, J. C. Muñoz, and J. C. Ferrer. Multiskilled workforce management by utilizing closed chains under uncertain demand: A retail industry case. *Computers & Industrial Engineering*, 127:74–88, 2019.

W. J. Hopp, E. Tekin, and M. P. Van Oyen. Benefits of skill chaining in serial production lines with cross-trained workers. *Management Science*, 50(1):83–98, 2004.

R. R. Inman, W. C. Jordan, and D. E. Blumenfeld. Chained cross-training of assembly line workers. *International Journal of Production Research*, 42(10):1899–1910, 2004.

R. R. Inman, D. E. Blumenfeld, and A. Ko. Cross-training hospital nurses to reduce staffing costs. *Health Care Management Review*, 30(2):116–125, 2005.

S. M. Iravani, M. P. Van Oyen, and K. T. Sims. Structural flexibility: A new perspective on the design of manufacturing and service operations. *Management Science*, 51(2):151–166, 2005a.

S. M. Iravani, M. P. Van Oyen, and K. T. Sims. Structural flexibility: A new perspective on the design of manufacturing and service operations. *Management Science*, 51(2):151–166, 2005b.

W. C. Jordan and S. C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4):577–594, 1995a.

W. C. Jordan and S. C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4):577–594, 1995b.

W. C. Jordan, R. R. Inman, and D. E. Blumenfeld. Chained cross-training of workers for robust performance. *IIE Transactions*, 36(10):953–967, 2004.

L. L. X. Li and B. E. King. A healthcare staff decision model considering the effects of staff cross-training. *Health Care Management Science*, 2(1):53–61, 1999.

M. Mac-Vicar, J. C. Ferrer, J. C. Muñoz, and C. A. Henao. Real-time recovering strategies on personnel scheduling in the retail industry. *Computers & Industrial Engineering*, 113:589–601, 2017.

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

T. Messelis, P. De Causmaecker, and G. Vanden Berghe. Algorithm performance prediction for nurse rostering. *Mista 2013*, pages 21–38, 2013. URL http://www.informatik.uni-osnabrueck.de/knust/class/.

OECD. *Health at a Glance 2013: OECD Indicators*. OECD Publishing, 2013. ISBN 9789264200715. doi: 10.1787/health_glance-2013-en.

OECD. *Health at a Glance 2019*. 2019. doi: https://doi.org/https://doi.org/10.1787/4dd50c09-en. URL https://www.oecd-ilibrary.org/content/publication/4dd50c09-en.

J. Olivella and D. Nembhard. Calibrating cross-training to meet demand mix variation and employee absence. *European Journal of Operational Research*, 248(2):462–472, 2016.

J. A. Oulton. The global nursing shortage: An overview of issues and actions. *Policy, Politics, and Nursing Practice*, 7(SUPPL/ 3):34–39, 2006. ISSN 15271544. doi: 10.1177/1527154406293968.

J. A. Paul and L. MacDonald. Modeling the benefits of cross-training to address the nursing shortage. *International Journal of Production Economics*, 150:83–95, 2014.

E. J. Pinker and R. A. Shumsky. The efficiency-quality trade-off of cross-trained workers. *Manufacturing & Service Operations Management*, 2(1):32–48, 2000.

R. Qin, D. A. Nembhard, and W. L. Barnes II. Workforce flexibility in operations management. *Surveys in Operations Research and Management Science*, 20(1):19–33, 2015.

E. Rahimian, K. Akartunalı, and J. Levine. A hybrid integer and constraint programming approach to solve nurse rostering problems. *Computers and Operations Research*, 82:83–94, 2017. ISSN 03050548. doi: 10.1016/j.cor.2017.01.016. URL http://dx.doi.org/10.1016/j.cor.2017.01.016.

D. Simchi-Levi and Y. Wei. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations research*, 60(5):1125–1141, 2012.

I. P. Solos, I. X. Tassopoulos, and G. N. Beligiannis. A generic two-phase stochastic variable neighborhood approach for effectively solving the nurse rostering problem. *Algorithms*, 6(2): 278–308, 2013.

G. K. Taskiran and X. Zhang. Mathematical models and solution approach for cross-training staff scheduling at call centers. *Computers & Operations Research*, 87:258–269, 2017.

V. M. Trivedi and D. M. Warner. A branch and bound algorithm for optimum allocation of float nurses. *Management Science*, 22(9):972–981, 1976.

G. Vairaktarakis and J. K. Winch. Worker cross-training in paced assembly lines. *Manufacturing & Service Operations Management*, 1(2):112–131, 1999.

J. Van den Bergh, J. Beliën, P. De Bruecker, E. Demeulemeester, and L. De Boeck. Personnel scheduling: A literature review. *European journal of operational research*, 226(3):367–385, 2013.

P. D. Wright and K. M. Bretthauer. Strategies for addressing the nursing shortage: Coordinated decision making and workforce flexibility. *Decision Sciences*, 41(2):373–401, 2010.

P. D. Wright and S. Mahar. Centralized nurse scheduling to simultaneously improve schedule cost and nurse satisfaction. *Omega*, 41(6):1042–1052, 2013.

K. K. Yang. A comparison of cross-training policies in different job shops. *International Journal of Production Research*, 45(6):1279–1295, 2007.