Framework for negotiation in Distributed Nurse Rostering Problems

Stefaan Haspeslagh[1], Patrick De Causmaecker[1], and Greet Vanden Berghe[2]

[1] K.U.Leuven Campus Kortrijk, Department of Computer Science,
E. Sabbelaan 53, 8500 Kortrijk, Belgium
{Stefaan.Haspeslagh,Patrick.DeCausmaecker@kuleuven-kortrijk.be}

[2] KaHo St.-Lieven, Information Technology,
Gebr. Desmetstraat 1, 9000 Gent, Belgium
Greet.VandenBerghe@kahosl.be

Abstract. This contribution deals with the distributed version of the nurse rostering problem. It is considered in a hospital with many separated wards. The nurse rostering problem within a ward is 'the local problem', rosters within a ward are 'local rosters'. At any time in the process, i.e. at the time of the construction of a roster for a certain time period as well as in the course of this time period when unexpected events cause rescheduling, wards may call in the help of their peers. At this level negotiation will take place. The details of the local rosters do not necessarily enter this negotiation level, and when they do, they may be translated from a local representation to a generally accepted vocabulary. After a motivation for a distributed approach, a general architecture is proposed and a negotiation protocol is described.

1 Introduction

The Nurse Rostering Problem (NRP) has been the subject of intensive study. The literature is very extensive and an overview is presented by Burke et al. [1]. The NRP is by nature a highly constrained problem and very hard to solve. Much attention has been paid to the local problem, i.e. at ward level. Rosters typically have to satisfy each nurse's work regulation while realizing a given coverage. Nurse preferences must be satisfied if possible. Apart from generating rosters for a fixed period of time, sudden staff shortages due to absence or unexpected overload must be coped with.

In the presented approach, we allow for requests to be raised by a ward in case of shortage. Other wards consider re-rostering in order to satisfy this demand. They submit weighted proposals to the requesting ward, the weight representing a cost associated with the rearrangement. When the specific request cannot be resolved, alternative proposals can be formulated. For instance, suppose that ward x signals a shortage of a nurse with a certain skill level within a particular shift and suppose that ward y cannot satisfy this request. Ward y can decide in that particular case to offer a nurse for another shift or with an equivalent skill

E. K. Burke, H. Rudová (Eds.): PATAT 2006, pp. 426–431. ISBN 80-210-3726-1.

level. Similarly, a request may specify a number of alternatives. We discuss this approach in more detail below.
Section 2 presents a brief overview of related research. Real world requirements related with the Nurse Rostering Problem are given in section 3. Using those requirements, section 4 argues why a distributed approach is appropriate. Section 5 investigates the criteria used to evaluate the rosters at each level. In section 6, we propose the ‘General Architecture’ designed to tackle the problem. Section 7 elaborates on an initial implementation. Finally, plans for future research are discussed.

2 Relation to published work

The nurse rostering problem that involves staff shortage, is often solved by re-planning under temporarily relaxed constraints. Several other solutions have been studied. Trivedi and Warner [4], for example, introduce a pool of *float nurses*. Siferd and Benton [5] resolve personnel shortage by calling nurses from other wards and by allowing nurses to work overtime.
At the level of the hospital, the size and the complexity of the problem become much larger. However, moving up to this level is not only a matter of scale. Some concerns disappear and others come up. Each ward, e.g., will typically maintain its own interpretation as to which extent certain constraints - not only preferences but work regulations as well - must be satisfied. Trying to solve the NRP at this level with the methods described in [1] is much harder if feasible at all. In the present contribution, we study the application of negotiation to tackle this problem.
This approach was pioneered by Kaplansky and Meisels in [6]. They introduce three stages. Scheduling agents generate a local solution in the first stage. The second stage constructs a globally feasible solution. The last stage is used by the scheduling agents to improve their local solutions. Kaplansky and Meisels assume a fixed pool of shared resources. Within our framework, all nurses are available for exchange.
Bard and Purnomo [7] introduce several ways to cope with personnel shortage. We mention those that are relevant for the present discussion. Initially, negotiation takes place with a pool of float nurses. Consequently, when a ward has more nurses than needed, the remaining nurses are placed on an on-call list. Other wards can call for those nurses when personnel shortage occurs. Finally, agency nurses are used. Replanning never takes place.

3 Real world requirements

The purpose of this section is to state the real world requirements in a hospital. In the next section, we use these requirements to show that they are naturally met by a distributed system.
A first class of requirements have to do with the scale of a hospital wide system and its associated maintainability problem. The second class of requirements

stem from the kind of information that is available either at the local or the global level and the inherent communication and interaction problems at the different levels.

The scale of a ward allows to build systems that can be used as real-time rostering systems [2]. This allows to react promptly to unexpected unavailabilities or temporary work overload. The manager can use the system for on-line consultancy and immediately discuss the possibilities with his staff. While hospital wide systems can be built [3] the execution times and the flexibility of the interface cannot be expected to allow for this kind of what-if analysis unless provisions are made to do one-ward-only rescheduling. This is closely related to the maintenance problem when moves of personnel cause changes in the local constraints and preferences.

The kind of information being processed at a local ward does not always have relevance or meaning at the hospital level. A specific nurse may not like to work on Thursdays or for interpersonal reasons two nurses should not make a team too often. This gives rise to preferences that a ward manager will want to take into account without spreading it throughout the hospital. A similar requirement goes for the work regulations. Although official individual work regulations are recorded at a central level, it might not make sense to have a central rostering system dealing with those regulations when building a hospital wide work roster. One rather wants these details to remain encapsulated at the ward level while reasoning about hospital wide concerns at a separate level. Among these concerns may be a fair distribution of resources and work load or strategic intention to favor a certain ward in order to increase a quality level or to handle a chronic safety problem.

We can thus conclude that a hospital nurse rostering system has to take at least two levels of decision making into account and that the criteria and objectives at the levels are essentially different. We argue in the next section that this requirement is naturally met by a distributed rostering system.

4 Motivation

In this section we describe the properties and benefits of a distributed method. First of all, a distributed approach offers a good *scalability*. When a new ward is organized, little effort must be made to support the new ward. Second, such a system is *flexible* and in a sense *easy to implement*. It can be introduced step by step. A couple of wards could initially implement the system. The system can then be gradually expanded with the other wards, but not every ward has to join. The way rostering is performed locally may differ at each ward. This way a large *autonomy* is provided to the wards. Our rostering approach also allows a large amount of communication between the wards. Since decisions are communicated, they might be more easily accepted.

5 Evaluation criteria

Following the two level example from section 3 we investigate the evaluation criteria and objectives that will be valid at each of the levels.
At the level of an individual ward, the criteria are those that have been analyzed in previous studies [2]. Broadly those are the satisfaction of the demands and the constraints with (individual) work regulations. Demands stem from the workload per time unit or shift in the ward and are typically expressed as a number of nurses of a specific qualification. Work regulations define, sometimes nurse-specific, values for e.g. the maximum number of consecutive night shifts, the minimum number of consecutive days off, and so on . . . Sometimes patterns are used to specify preferred work time distributions. These constraints are often very complicated and their modeling is a highly demanding task. In addition one wants a system that allows to specify personal preferences and ad hoc requests.We revisit implementation issues in the light of the present study during the presentation.
At the hospital level other criteria must be taken into account. At this level the manager wants to guard the fairness of the work load distribution, wants to raise certain quality levels, wants to decide where resource shortages are allowable at peak moments and so on . . .
A model must allow to express the requirements at all levels. In this contribution we will discuss a negotiation based model. Wards will be represented by autonomous agents. They will generate rosters along well investigated lines, but they will be able to detect a local shortage and to formulate requests to solve shortages. The agents will also be able to interpret requests and generate answers to requests. Our negotiation protocol is designed in such a way that it allows to express local needs at the higher level, the hospital level. The decision criteria for individual wards to respond to requests and to accept proposals are used to express the hospital level requirements. We develop an example during the presentation.

6 General Architecture

The General Architecture consists of three building blocks. A first building block is a common language. This language enables the wards to formulate internal information into a representation that each ward understands. The second building block is a negotiation protocol. This protocol allows a ward to formulate multiple questions and to select one of the alternatives offered. The third building block represents a ward. Each ward (e.g. $W1$, $W2$ and $W3$) is an autonomous entity that consists of the following components:

- IDP: When a problem in a ward occurs, the Internal Problem Descriptor specifies the problem.
- EDP: The External Problem Descriptor translates an internal problem into a request in terms of a common language known by all wards.

- ERD: The External Request Descriptor transforms an incoming request into an internal representation.
- IRD: The Internal Request Descriptor passes the request to the internal engines.
- IAD: A proposal to meet the request is formulated into an answer by the Internal answer Descriptor
- EAD: The External Answer Descriptor translates the answer into a common language.
- EOD: An offer by a ward is translated into internal terms by the External Offer Descriptor
- IOD: The Internal Offer Descriptor handles offers.

They will be discussed in future work.

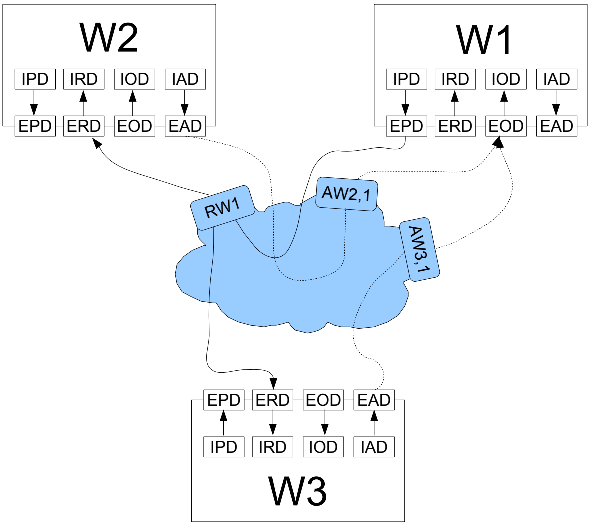

Fig. 1. General Architecture

An overview of the architecture is presented in Fig. 1. RW_x denotes a request by Ward x, and $AW_{y,x}$ denotes Ward y's answer to Ward x.

7 Initial solution

In a first iteration of this problem, two building blocks are worked out. The common language allows wards to raise specific requests in case of personnel shortage and to submit proposals to satisfy demands. Particular attention is paid to the assurance that minimal information needs to be exchanged for formulating requests and answers. Secondly, a negotiation protocol is designed. In order to test our approach, a ward is elaborated with simplifications of the P- and R-components.

8 Future research

Our initial common language supports full answers to requests. Further elaboration must allow for partial proposals to satisfy requests. A more sophisticated negotiation protocol needs to be worked out. Similarities and repetition of requests should be analyzed. E.g. if ward x always lacks a nurse for a certain shift, and ward y is always able to respond to this demand, a transfer of personnel from y to x may be considered.
Intelligence must be built into the components introduced in Section 6. For example, an internal request by the EPD could be translated in different versions destined for different wards. The EPD can *learn* how a particular ward responds to certain requests and attempt to formulate the requests in such a way so that the ward will certainly respond.

References

1. E.K. Burke, P. De Causmaecker, G. Vanden Berghe, H. Van Landeghem. The State of the Art of Nurse Rostering. In *Journal of Scheduling*, 2004, Vol. 7, No. 6, 441-499 (2004)
2. E.K. Burke, P. De Causmaecker, G. Vanden Berghe: Novel Meta-heuristic Approaches to Nurse Rostering Problems in Belgian Hospitals, Chapter 44 in J. Leung: Handbook of Scheduling: Algorithms, Models and Performance Analysis, CRC Press, 2004, 44.1-44.18
3. E.K. Burke, P. De Causmaecker, S. Petrovic, G. Vanden Berghe: Metaheuristics for handling Time Interval Coverage Constraints in Nurse Scheduling, Applied Artificial Intelligence, Vol. 20, No. 3, 2006 (to appear)
4. Trivedi V.M. and M. Warner. A branch and bound algorithm for optimum allocation of float nurses. In *Management Science*, *22*(9), 972-981 (1976)
5. Siferd S.P. and W.C. Benton. Workforce staffing and scheduling: Hospital nursing specific models. In *European Journal of Operational Research*, *60*, 233-246 (1992)
6. E. Kaplansky and A. Meisels. Distributed Personnel Scheduling - Negotiation among Scheduling Agents. In *Annals of Operations Research*, (to appear)
7. J.F. Bard and H.W. Purnomo. Real-Time Scheduling for Nurses in Response to Demand Fluctuations and Personnel Shortages. M.A. Trick and E.K. Burke (editors) Proceedings of the 5th International Conference on the Practice and Theory of Automated Timetabling, Pittsburgh, 67-87 (2004)