

Distributed Constraint Optimization for Medical Appointment Scheduling

Markus Hannebauer and Sebastian Müller

GMD — German National Research Center for Information Technology
Research Institute for Computer Architecture and Software Technology (FIRST), Berlin
markus.hannebauerl@sebastian.mueller@gmd.de

1. INTRODUCTION

Scheduling problems (that is assigning resources and time points to given tasks) arise in many real-world domains. Recently, several researchers have applied agent technology to such problems especially in manufacturing and business process control [4]. Because of the widely spread lack of a theoretical foundation for this kind of collaborative problem solving, the constraint satisfaction/optimization community has recently extended the theoretically well-founded *constraint satisfaction problem* (CSP) model to a distributed setting, called *distributed constraint satisfaction problems* (DCSPs) [6]. Research in DCSP solving has in most cases directly succeeded CSP research considering mainly academic problems with binary constraints and complete algorithms for very small instances.

Medical appointment scheduling can be modelled as such a DCSP, but not very conveniently using only binary constraints. Being related to meeting scheduling [3], rostering [5] and patient scheduling [1], medical appointment scheduling involves several complex constraints only known in this domain. Therefore, we propose to deploy several CSP solvers that can handle complex constraints and coordinate them. First, we will briefly show how to model the medical appointment scheduling problem for such an approach, then how to coordinate several solvers to cope with it.

2. PROBLEM MODEL

DEFINITION 1. A distributed constraint optimization problem (DCOP) is specified by $\Pi = (X, D, C, \phi)$.

- $X = \{x_1, \dots, x_n\}$ is a set of variables x_i each ranging over its domain D_i from the set $D = \{D_1, \dots, D_n\}$. A labeling $\lambda = (v_1, \dots, v_n) \in D_1 \times \dots \times D_n$ assigns a value $v_i \in D_i$ to each variable x_i .
- $C = \{C_1, \dots, C_m\}$ is a set of constraints $C_i \subseteq D_{i_1} \times \dots \times D_{i_k}$ restricting the set of feasible labelings.
- $\phi = \{(a_1, X_1 \cup C_1, o_1), \dots, (a_p, X_p \cup C_p, o_p)\}$ is a par-

tion of $X \cup C$ and assigns an agent a_i with an optimization criterion o_i to each partition block.

Given this specification, the problem is to find a feasible labeling λ^* that has the highest combined ranking by all agents. I.e. given a max operator that determines the highest ranked multi-objective vector among $(o_1(\lambda), \dots, o_p(\lambda))$ for all λ , we are searching for $\lambda^* \in \{(v_1, \dots, v_n) \in D_1 \times \dots \times D_n \mid \forall C_i \in C : (v_{i_1}, \dots, v_{i_k}) \in C_i\} = \Lambda$ such that $(o_1(\lambda^*), \dots, o_p(\lambda^*)) = \max\{(o_1(\lambda), \dots, o_p(\lambda)) \mid \lambda \in \Lambda\}$.

The difficulty in applying this mathematical model to the practice of medical appointment scheduling lies in the identification of X , D , C and ϕ in the scenario. We have done this following a concrete case study at Charité Berlin, Europe's largest hospital. The variables of scheduling problems are the time point and resource assignments of tasks and the domains are given by all possible time points and all available resources. We are using a discrete and non-branching model of time to build a finite scheduling horizon T . In our case study, resources correspond to *diagnostic units* comprising staff and several *workplaces*. Workplaces provide a set of appointment types and are only available for parts of the horizon. *Patients* request a partially ordered set of *appointments* that each specify an appointment type as well as a demand for certain staff. The link between appointments and time/resources is modeled by a set of *assignments* $X = \{(t_{\text{start}}, \Delta t_{\text{dur}}, w)\}$ that represent a starting time point, a duration and a workplace from the set of all workplaces W for each appointment. Hence, the set of assignments is the set of variables in our DCOP model and the domains are all defined by a three-dimensional space $D_1 = \dots = D_n = (T \times T \setminus \{0\} \times W)$.

The constraints C in medical appointment scheduling are imposed by patients (C_{pat}) and diagnostic units (C_{unit}). Patients enter the clinic on a certain time point and are not available at nights. Additionally, they define a medically motivated partial order among their appointments and cannot attend more than one appointment at once. Workplaces of diagnostic units have to provide the types of the appointments assigned to them, they have to be available, need certain setup times between two appointments and can also participate only in one appointment at a time. The staff assigned to the diagnostic units demands a diversified work profile to keep them from monotonic work and restricts the number of workplaces in a diagnostic unit that can run in parallel.

To define the partition ϕ we assign patients, their appointments and constraints to *patient agents* and diagnostic units, their workplaces and constraints to *diagnostic unit agents*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AGENTS'01, May 28-June 1, 2001, Montréal, Quebec, Canada.

Copyright 2001 ACM 1-58113-326-X/01/0005 ...\$5.00.

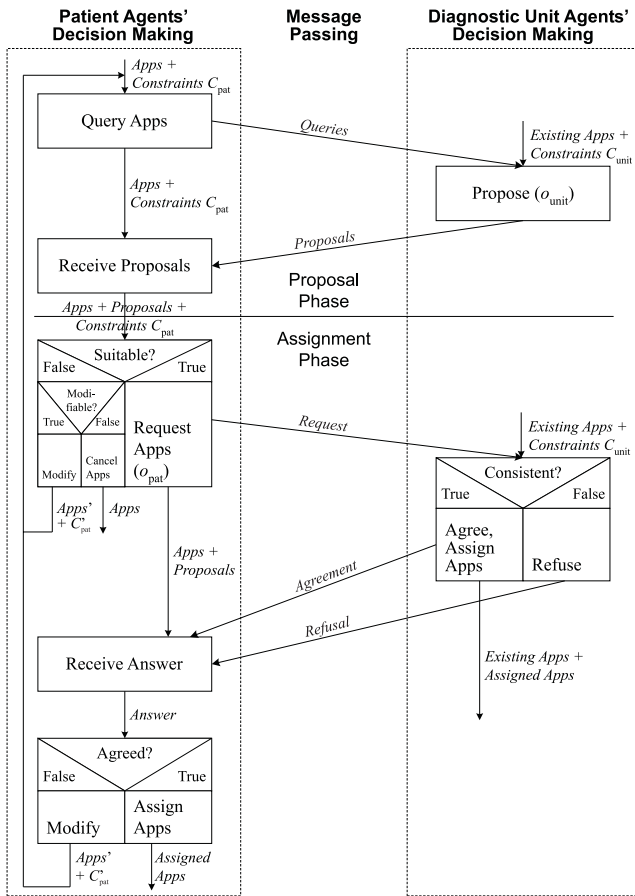


Figure 1: Multi-phase Agreement Finding

Patient agents use an optimization criterion o_{pat} that favors a tight appointment sequence by minimizing the derivation from the mean starting time of all appointments of a patient. Diagnostic unit agents maximize the calendar density (o_{unit}), i.e. the quotient of assigned time and availability, and such the usage of their resources.

3. SOLVER COORDINATION

For the coordination of the several problem solving agents we propose the *multi-phase agreement finding* (MPAF) protocol.¹ MPAF is a distributed algorithm that starts with an empty labeling and incrementally constructs a feasible labeling by adding consistent variable assignments. It consists of several phases that comprise asynchronous constraint propagation, asynchronous search and retraction mechanisms. For medical appointment scheduling we have adapted this algorithm to a committed choice algorithm without backtracking. Hence, it is not complete and may miss some feasible solutions. On the other hand, it allows for constraint relaxation.

The MPAF instance for appointment scheduling involves a set of patient agents and a set of diagnostic unit agents. Patient agents collect constraint information on their appointments and propose assignments according to this. If the diagnostic unit agents agree to the assignments, an agreement has been found and the patient agent as well as the diagnostic unit agents will commit to the agreed assignments.

¹A more formal description of MPAF can be found in [2].

Figure 1 provides a flow model of MPAF. The idea of the *proposal phase* is to anticipate future conflicts and such reduce search effort by making decisions on variable labelings well-founded. A patient agent starts this phase by selecting a subset *Apps* from the patient's open appointments, considering its own constraints C_{pat} and sending according query messages to the set of all potential providing diagnostic unit agents. The diagnostic unit agents answer the queries according to their knowledge about their constraints C_{unit} and already existing appointments. Doing so, unit agents use their own optimization criteria o_{unit} to restrict their proposals to good ones. The patient agent waits a certain time for these proposals and collects them.

In the *labeling phase* the patient agent decides whether the received proposals are suitable or not. If not, it tries to modify the problem (e.g. by relaxing C_{pat}). Then it restarts the protocol. If the proposals are suitable, the patient agent will select the best among them following o_{pat} and send requests to a newly formed subset of diagnostic unit agents. Because of asynchronous change, the proposal of a certain diagnostic unit agent may not hold anymore. Hence, diagnostic unit agents have to check C_{unit} and the existing appointments again. If the requested assignment is consistent, the assignment is added to the existing ones on diagnostic unit agent side and an agree message is sent. If not, a refuse message is sent and nothing changes on diagnostic unit agent side. The patient agent receives an agreement or refusal to a subset of its requested appointments and reacts accordingly by assigning the agreed appointments or modifying refused appointments and restarting the protocol.

4. IMPLEMENTATION

To assess our approach we are currently implementing the ideas presented here. Building upon standard middleware such as DCOM, we have created agents that use SICStus Prolog to make their decisions in the MPAF protocol. We have also realized an on-line simulator that creates a random but realistic medical appointment scheduling scenario. Patient agents and diagnostic unit agents store their scheduling results in a database. This data can then be used to determine quality and performance measures, such as mean make span of patients, calendar density of patients and diagnostic units and patient throughput. Performance measures are the number of exchanged messages and CPU time. Results from this empirical study are expected to appear soon.

5. REFERENCES

- [1] K. Decker and J. Li. Coordinated hospital patient scheduling. *Proc. ICMA-98*, Paris, France, 1998.
- [2] M. Hannebauer. How to model and verify concurrent algorithms for distributed CSPs. *Proc. CP-2000, LNCS* 1894. Springer, 2000.
- [3] S. Sen. An automated distributed meeting scheduler. *IEEE EXPERT*, 12(4), 1997.
- [4] W. Shen and D. Norrie. Agent-based systems for intelligent manufacturing: A state-of-the-art survey. *KNOWL INFORM SYST*, 1(2), 1999.
- [5] G. Solovetsky and E. Gudes. Solving a real-life time tabling and transportation problem using distributed CSP techniques. *Proc. CP-96 WS Constraint Applications*, Cambridge, USA, 1996.
- [6] M. Yokoo, E. Durfee, T. Ishida, and K. Kuwabara. The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE T KNOWL DATA EN*, 10(5), 1998.