Learning the Parameters of Global Constraints for Medical Scheduling

Émilie Picard-Cantin[1], Mathieu Bouchard[2],
Claude-Guy Quimper[1], and Jason Sweeney[2]

[1] Université Laval
emilie.picard-cantin.1@ulaval.ca, claude-guy.quimper@ift.ulaval.ca
[2] PetalMD
mathbouchard@gmail.com, jason.pierre.sweeney@gmail.com

Keywords: Constraint Acquisition, Timetabling, Machine Learning, CSP, Global Constraints, Brand-and-Bound

In medical scheduling, one wants to model and to compute a schedule that satisfies both the requirements and the availability of the medical staff. Modelling such a problem necessitates specific knowledge both in constraint optimization and in medical scheduling. Each medical team has its own requirements and specificities. This results in multiple models, one for each medical team. Often, some constraints of the problem are already known or obvious (ex: a person cannot work at two different locations simultaneously). But there exist known constraints for which the parameters differ from one medical team to another and that need to be encoded, for each model, by the modelling expert. For instance, should a person work a maximum 3 night shifts per period of 7 days or 4 night shifts per period of 9 days? Defining such constraints are time consuming for the modelling experts, especially when there are many.

In order to automate the modelling process, we propose to learn the parameters of the global constraints from schedules that were previously used. We show how to compute the parameters that are most likely the ones that generated a given set of schedules. The idea is to cleverly explore the possible combinations taken by the constraint's parameters without explicitly enumerating all combinations. We apply our method to learn parameters of global constraints used in timetabling problems such as Sequence and SubsetFocus. The later constraint is our adaptation of the constraint Focus to timetabling problems.