Time-Bucket Relaxation Based Mixed Integer Programming Models for Scheduling Problems: A Promising Starting Point for Matheuristics*

Günther R. Raidl, Thomas Jatschka, Martin Riedler, and Johannes Maschler

Institute of Computer Graphics and Algorithms,
TU Wien, Vienna, Austria
{raidl|riedler|maschler}@ac.tuwien.ac.at, jatschka.thomas@gmail.com

1 Introduction

In job shop and project scheduling problems, generally speaking, a set of activities shall be scheduled over time. The execution of the activities typically depends on certain resources of limited availability and diverse other restrictions like precedence constraints. A feasible schedule is sought that minimizes some objective function like the makespan. For such problems, *mixed integer linear programming* (MIP) techniques are frequently considered, but also known to have severe limitations.

Basically, there are few general MIP modeling strategies for approaching such scheduling problems: Firstly, it is sometimes possible to come up with a compact model where the starting times of activities are directly expressed by means of corresponding variables. Resource constraints, however, impose a particular challenge in this respect. While they can be often treated in principle, e.g., by discrete-event models [2], these models are typically rather weak. A second, frequently applied option are so-called *time-indexed* (TI) formulations. They are based on a discretization of time, i.e., the activities may only start on a limited set of possible starting times. Binary variables are used that are additionally indexed by these possible starting times. The success of such TI models strongly depends on the resolution of the time discretization. While such models can have strong *linear programming* (LP) relaxations, the number of variables and constraints increases dramatically with the number of possible starting times. Frequently, a rather crude discretization can therefore only be applied to obtain any result in reasonable computation time. Further MIP techniques for approaching the considered scheduling problems make use of exponentially sized models and apply advanced techniques such as column generation, Lagrangian decomposition, or Benders decomposition, see, e.g., [2]. While they are frequently very successful, they are also substantially more complex to develop and implement.

Here, we consider a relaxation of a potentially very fine-grained TI model in which the set of possible starting times is partitioned into so-called *time-buckets*

* We thank EBG MedAustron GmbH, Wiener Neustadt, Austria, for the collaboration on particle therapy patient scheduling and partially funding this work.

(TB). This TB relaxation is typically much smaller than the original TI model and can be solved relatively quickly. An obtained solution provides a lower bound for the TI model's solution value but in general does not directly represent a feasible schedule as activity start times are only restricted to certain time-intervals. This solution, however, provides a promising starting point for matheuristics. On the one hand, we may try to derive a feasible schedule by heuristically fixing the start times to specific values, trying to fulfill all constraints. On the other hand, we can further subdivide some time-buckets and re-solve the resulting refined model to obtain an improved bound and a model that comes closer to the TI model. Doing this refinement iteratively yields a matheuristic that in principle converges to a provably optimal solution. In practice, it is crucial to subdivide the time-buckets in a sensible way in order to increase the model's size only slowly while hopefully obtaining significantly stronger bounds. (Meta-)heuristic techniques and dual variable information may provide a strong guidance.

The basic idea of the time-bucket relaxation originates in work from Wang and Regan [4] on the traveling salesman problem with time windows. Dash et al. [1] build upon this work and suggest an iterative refinement based on the solution to the LP-relaxation. We are not aware of any work that applies this principle already in the scheduling domain. There is just other work where the TI model is applied with different resolutions for the time discretization, but such approaches do in general not yield lower bounds and introduce imprecisions and are therefore conceptually different.

2 Resource Constrained Scheduling with Precedence Constraints

The above sketched general approach is more specifically investigated on a resource constrained scheduling problem with precedence constraints. This problem, for example, arises as a subproblem in the daily planning of activities to treat cancer patients with modern particle therapy [3]. Our experimental evaluation considers benchmark instances from this application.

We are given a set of resources $R = \{1, \ldots, \rho\}$, a set of activities $A = \{1, \ldots, \alpha\}$, and for each activity $a \in A$ a processing time p_a, a release time t^{r}_a, a deadline t^{d}_a, and a subset of required resources $Q_a \subseteq R$. Let the overall (huge) set of discrete times be $T = \{T^{\min}, \ldots, T^{\max}\}$. Each resource $r \in R$ is only available at certain time intervals specified by set $W_r \subseteq T$. Last but not least, precedence constraints among the activities are stated by a directed acyclic graph $G = (A, P)$ with $P \subset A \times A$ and for each precedence relation expressed by an arc $(a, a') \in P$ minimum and maximum end-to-start time lags $L^{\min}_{a,a'}, L^{\max}_{a,a'} \in \mathbb{N}_{\geq 0}$ with $L^{\min}_{a,a'} \leq L^{\max}_{a,a'}$ need to be obeyed.

A solution $S = (S_1, \ldots, S_\alpha) \in T^\alpha$ assigns to each activity $a \in A$ a starting time $S_a \in T$, from which on the activity is performed without preemption. We are looking for a feasible solution that minimizes the makespan.

3 Time-Bucket Relaxation and Matheuristic

Let $B = \{B_1, \dots, B_\beta\}$ be a partitioning of T into subsequent time-buckets with $B_b = \{B_b^{\text{start}}, \dots, B_b^{\text{end}}\}, \forall b = 1, \dots, \beta$, and $B_b^{\text{end}} + 1 = B_{b+1}^{\text{start}},\ \forall b = 1, \dots, \beta - 1$. We further make the following definitions.

- $I(B) = \{1, \dots, \beta\}$ is the index set referring to all buckets in B.
- $W_r^B(b) = |B_b \cap W_r|$ denotes the aggregated availability of resource $r \in R$ over the whole bucket $b \in I(B)$.
- $C_a = \{C_{a,1}, \dots, C_{a,\gamma_a}\} \subseteq 2^{I(B)}$ refers to all subsets of consecutive buckets in B to which an activity $a \in A$ can be jointly assigned so that some part of activity a is performed in each of the buckets. These sets can be determined by "sliding" the activity over all time-slots and taking the covered buckets.
- Let $t_{a,c}^{\text{smin}}$ be the earliest time-slots from T at which activity a can possibly start when it is assigned to bucket sequence $C_{a,c}$, and $t_{a,c}^{\text{smax}}$ the latest.
- For each bucket sequence $C_{a,c} \in C_a$ and each contained bucket $b \in C_{a,c}$ we further determine a lower bound $z_{a,b,c}^{\min}$ and an upper bound $z_{a,b,c}^{\max}$ for the number of time-slots at which activity a can possibly take place in bucket b when activity a is assigned to $C_{a,c}$.

The TB relaxation can now be stated as follows.

$$\min\ MS \tag{1}$$

$$\sum_{c=1}^{\gamma_a} y_{a,c} = 1 \qquad \forall a \in A \tag{2}$$

$$\sum_{c=1}^{\gamma_a} t_{a,c}^{\text{smin}} \cdot y_{a,c} + p_a \leq MS \qquad \forall a \in A \tag{3}$$

$$\sum_{a \in A, C_{a,c} \in C_a \mid b \in C_{a,c} \wedge r \in Q_a} z_{a,b,c}^{\min} \cdot y_{a,c} \leq W_r^B(b) \qquad \forall r \in R,\ b \in I(B) \tag{4}$$

$$\sum_{c'=1}^{\gamma_{a'}} t_{a',c'}^{\text{smax}} \cdot y_{a',c'} - \sum_{c=1}^{\gamma_a} t_{a,c}^{\text{smin}} \cdot y_{a,c} \geq p_a + L_{a,a'}^{\min} \qquad \forall (a,a') \in P \tag{5}$$

$$\sum_{c'=1}^{\gamma_{a'}} t_{a',c'}^{\text{smin}} \cdot y_{a',c'} - \sum_{c=1}^{\gamma_a} t_{a,c}^{\text{smax}} \cdot y_{a,c} \leq p_a + L_{a,a'}^{\max} \qquad \forall (a,a') \in P \tag{6}$$

$$y_{a,c} \in \{0,1\} \qquad \forall a \in A, c = 1, \dots, \gamma_a \tag{7}$$

$$MS \geq 0 \tag{8}$$

Variable MS represents the makespan to be minimized (1). Binary variables $y_{a,c}$ indicate if activity $a \in A$ is completely performed in bucket sequence $C_{a,c}$. Equations (2) ensure that for each activity exactly one bucket sequence is chosen from C_a. Inequalities (3) are used for determining the makespan MS. Inequalities (4) consider for each time bucket the aggregated resource availabilities and resource consumptions for performing the respective activities. Finally, inequalities (5) and (6) represent the precedence constraints with the minimum and maximum time lags, respectively.

Our matheuristic works as follows. We initially solve the TB relaxation for a rather crude partitioning of T into buckets. Then we try to derive a feasible schedule from the solution of the TB relaxation, i.e., we try to choose valid activity starting times as far as possible in correspondence to the selected bucket sequences. This is done by a greedy construction heuristic that considers the time buckets in chronological order and the assigned activities in a topological order, taking care of the precedence constraints and resource constraints as far as possible. Should we be able to find a feasible schedule whose makespan corresponds to the solution value of the TB relaxation, then this schedule is optimal and we can terminate.

Otherwise, the bucket partitioning is further refined by splitting buckets related to violated constraints. Furthermore, valid inequalities cutting off current infeasibilities may be added to the model. The refined model is solved again and the whole process iterated. We investigate and compare several strategies for the bucket splitting, considering also dual variable information from the relaxation.

4 Results and Conclusions

An experimental comparison with a compact discrete-event model and a classical TI formulation clearly shows the advantages of the TB relaxation: While the discrete-event model is only applicable to tiny instances due to its poor LP relaxation, the TI formulation suffers from its huge size when considering practically reasonable time discretizations. The matheuristic based on the iterative refinement of the TB model, however, soon yields reasonable lower bounds as well as feasible heuristic solutions, and both are improved over time.

The described approach is relatively generic and can rather easily be adapted to related scheduling problems. Clearly, there are many ways to enhance the basic concept: Intermediate heuristic solutions may be further improved by advanced local search techniques, the model may be strengthened by additional valid inequalities, possibly extending the approach to a branch-and-cut algorithm. More generally the field of hybrid metaheuristics and matheuristics provides plenty of opportunities to further exploit the proposed time-bucket relaxation.

References

1. Dash, S., Günlük, O., Lodi, A., Tramontani, A.: A time bucket formulation for the traveling salesman problem with time windows. INFORMS Journal on Computing 24(1), 132–147 (2012)
2. Hooker, J.N.: Planning and scheduling by logic-based Benders decomposition. Operations Research 55(3), 588–602 (2007)
3. Maschler, J., Riedler, M., Stock, M., Raidl, G.R.: Particle therapy patient scheduling: First heuristic approaches. In: Proceedings of the 11th Int. Conference on the Practice and Theory of Automated Timetabling (to appear 2016)
4. Wang, X., Regan, A.C.: On the convergence of a new time window discretization method for the traveling salesman problem with time window constraints. Computers & Industrial Engineering 56(1), 161–164 (2009)