

Curriculum based course timetabling: new solutions to Udine benchmark instances

Gerald Lach · Marco E. Lübbecke

© Springer Science+Business Media, LLC 2010

Abstract We present an integer programming approach to the university course timetabling problem, in which weekly lectures have to be scheduled and assigned to rooms. Students' curricula impose restrictions as to which courses may not be scheduled in parallel. Besides some hard constraints (no two courses in the same room at the same time, etc.), there are several soft constraints in practice which give a convenient structure to timetables; these should be met as well as possible.

We report on solving benchmark instances from the literature and the 2nd International Timetabling Competition which are based on real data from the university of Udine. The first set is solved to proven optimality; for the second set we give solutions which on average compete well with or beat the previously best known solutions. Our algorithm is not an overall winner, but it is very robust in the sense that it deterministically gives satisfactory lower and upper bounds in reasonable computation time without particular tuning. For slightly larger instances from the literature our approach shows significant potential as it considerably beats previous benchmarks. We further present solutions of proven quality to a few much larger instances with more elaborate hard constraints.

Keywords Integer programming · Decomposition · University course timetabling

1 Introduction

The curriculum based course timetabling problem is to assign weekly lectures to time periods and rooms in such a way that a number of obvious hard constraints are fulfilled: a

G. Lach (✉)
Institut für Mathematik, Technische Universität Berlin, MA 7-2, Straße des 17. Juni 136, 10623 Berlin,
Germany
e-mail: lach@math.tu-berlin.de

M.E. Lübbecke
Fachbereich Mathematik, AG Optimierung, Technische Universität Darmstadt, Dolivostr. 15,
64293 Darmstadt, Germany
e-mail: luebbecke@mathematik.tu-darmstadt.de

teacher can only teach one course at a time, a lecture room cannot host two courses simultaneously, courses of the same curriculum must not be scheduled in parallel, etc. If this is impossible, the number of violations is to be minimized. Furthermore, several soft constraints should be met as well as possible; these typically give desired structural properties like coherent daily time slots for lectures of the same curriculum, etc. This problem, also known as (curriculum based) university course timetabling, received much attention in the operations research literature, see the surveys (Burke and Petrovic 2002; Schaerf 1999), not least due to the fact that practical data is available for benchmarking, in particular instances from the university of Udine (Di Gaspero and Schaerf 2003, 2006). In ITC2007, the second International Timetabling Competition (www.cs.qub.ac.uk/itc2007), more benchmarks from Udine were provided (Di Gaspero et al. 2007; McCollum et al. 2009), together with extended problem definitions, in particular for the soft constraints.

In this paper, we report for the first time on solving the four original Udine instances to proven optimality (which is also, but certainly not only due to the fact that they became rather easy for modern integer programming solvers), and give solutions to the 2007 instances which do not violate any hard constraint. Here it turns out that we are able to very well compete with, and often beat the strongest known, tailored solution methods which are based on heuristics. We furthermore provide solutions to instances derived from practical data from Berlin's Technical University which feature slightly more elaborate hard constraints.

We approach the problem (which is NP-hard) via integer programming, as has been proposed before, see e.g., Burke et al. (2007, 2008b, 2008), Carter (2001), Daskalaki and Birbas (2005), Daskalaki et al. (2004), Qualizza and Serafini (2005), Schimmelpfeng and Helber (2007). However, instead of directly solving a natural formulation based on three-indexed variables for the course/time/room assignment, we decompose the problem in two stages. In the first stage, we only match time periods and lectures; these pairs are then feasibly assigned to rooms in a second step. This decomposition is exact with respect to hard constraints, that is, no solutions are lost. This can be achieved by implicitly taking care about feasibility for room assignments already in the first stage. Overall, this approach results in easier integer programs, and thus larger instances can be solved.

University course timetabling

Each course consists of several lectures, each day consists of several time slots. A pair of day and time slot is called a period. A curriculum is a set of courses no two of which may be scheduled in parallel. Every lecture has to be scheduled to a period in a room which provides enough seats to host the lecture (in our Berlin instances, each room must also provide the requested features like projector, PC, blackboard, location, etc.). No two courses by the same teacher or which appear in the same curriculum may be scheduled at the same period; no two courses may take place at the same period in the same room. All these constraints are considered *hard*. We defer the subtleties of soft constraints to our discussion on how to formulate them in our integer programs in Sect. 3.

In a companion paper (Lach and Lübbecke 2008) we developed the theoretical background for our decomposition which considered hard constraints only. Here, we report on how to make it useful in practice, in particular, we state how to incorporate a variety of soft constraints.

2 The hard constraint solver framework

Our focus is on meeting all hard constraints (resp. as many as possible); thus, the core of our model is built around this goal. Soft constraints are added as needed; see Sect. 3.

Denote by \mathcal{C} the set of courses, by \mathcal{R} the set of rooms, and by \mathcal{P} the set of periods. For each course $c \in \mathcal{C}$ we know its eligible periods $P(c) \subseteq \mathcal{P}$, eligible rooms $R(c) \subseteq \mathcal{R}$, and that $\ell(c)$ lectures have to be scheduled; that is, we have to provide $\ell(c)$ different periods for course c . As an example objective function (we use in Berlin) we formulate teachers' preferences $prio(c, p)$ for course/period combinations; the smaller the number, the higher the priority.

Time conflicts of any kind are represented via a conflict graph $G_{\text{conf}} = (V_{\text{conf}}, E_{\text{conf}})$: A vertex (c, p) represents an eligible combination of course c and period p (and our model has a binary variable $x_{c,p} \in \{0, 1\}$ for each such combination). Two nodes (c_1, p_1) and (c_2, p_2) are adjacent iff it is forbidden that c_1 is scheduled at p_1 and c_2 is scheduled at p_2 (typically, $p_1 = p_2$). This stable set characteristic of the problem motivated several researchers to relate timetabling to graph coloring, see e.g., Burke et al. (2007), and references therein.

Instead of using binary variables which represent whether course c is scheduled at period p in room r , we reduce the problem in three dimensions to a problem in two dimensions, implicitly taking care of room conflicts. To this end, we represent eligible combinations of courses and rooms as undirected bipartite graphs $G_p = (\mathcal{C}_p \cup \mathcal{R}_p, E_p)$, one for every period $p \in \mathcal{P}$. Courses which may be scheduled at p are given in set \mathcal{C}_p , and \mathcal{R}_p denotes the set of all eligible rooms for all courses in \mathcal{C}_p . A course c and a room r are adjacent iff r is eligible for c . For ease of exposition let $G = (\mathcal{C} \cup \mathcal{R}, E)$ be the graph consisting of all components $G_p, p \in \mathcal{P}$.

For a subset $U \subseteq \mathcal{C}$ of vertices, denote by $\Gamma(U) := \{i \in \mathcal{R} \mid j \in U, (i, j) \in E\}$ the neighborhood of U . Hall's stable marriage theorem (Lovász and Plummer 1986) states that a bipartite graph $G = (\mathcal{C} \cup \mathcal{R}, E)$ has a matching of all vertices in \mathcal{C} into \mathcal{R} if and only if $|\Gamma(U)| \geq |U|$ for all $U \subseteq \mathcal{C}$. Enforcing this condition, we are able to schedule courses in such a way that rooms can be assigned later.

We call this the *first stage of the decomposition*. The resulting integer program obviously has an exponential number of constraints (3), and we give details in Lach and Lübbecke (2008) on how to cope with this (and why in practice there are not too many of them).

$$\min \quad \sum_{(c,p) \in V_{\text{conf}}} prio(c, p) \cdot x_{c,p} \tag{1}$$

$$\text{s.t.} \quad \sum_{p \in P(c)} x_{c,p} = \ell(c) \quad \forall c \in \mathcal{C}, \tag{2}$$

$$\sum_{c \in U} x_{c,p} \leq |\Gamma(U)| \quad \forall U \subseteq \mathcal{C}, p \in \mathcal{P}, \tag{3}$$

$$x_{c_1,p_1} + x_{c_2,p_2} \leq 1 \quad \forall ((c_1, p_1), (c_2, p_2)) \in E_{\text{conf}}, \tag{4}$$

$$x_{c,p} \in \{0, 1\} \quad \forall (c, p) \in V_{\text{conf}}. \tag{5}$$

Once this program is solved, the *second stage of the decomposition* merely consists of solving a sequence of minimum weight bipartite perfect matching problems in polynomial time, one for each period $p \in \mathcal{P}$. Clearly, this decomposition approach is exact, that is, in principle it deterministically finds optimal solutions, provided one allows enough running time.

3 Integrating soft constraints

Besides mandatory constraints there is a wealth of possibilities for constraints which cannot be kept in general, but best possibly fulfilling them gives desired structures to timetables. For these soft constraints, we stick to the definitions from ITC2007, see again Di Gaspero et al. (2007) and the more recent McCollum et al. (2009). Four types are mentioned (and defined below): *RoomCapacity (RC)*, *MinimumWorkingDays (MWD)*, *CurriculumCompactness (CC)*, and *RoomStability (RS)*. The first three can easily be included in the first stage of the decomposition. On the other hand, the RS constraints need to go in the second stage, and are ignored in the first. As a consequence, we theoretically may miss a globally optimal solution, even when both stages are optimally solved. However, in that case, solution quality would not significantly decrease since the RS constraints are the least important soft constraints. Penalties for violations are taken from Di Gaspero et al. (2007) and McCollum et al. (2009).

3.1 RoomCapacity constraints

A room should provide as many seats as requested by each assigned course. A penalty occurs for each missing seat. This constraint is a hard constraint in our original (Berlin) framework; ITC2007, however, treats it as soft. One might expect to handle room capacity in the second stage of the decomposition, but a modification of Hall's conditions (3) already does the job.

Let p be an arbitrary but fixed period. Hall's conditions (3) are replaced by the following set of constraints. We first require the number of courses that can take place at p to be at most the number of available rooms:

$$\sum_{c \in \mathcal{C}} x_{c,p} \leq |\mathcal{R}| \quad \forall p \in \mathcal{P}. \quad (6)$$

This avoids conflicts in the assignment of rooms. Next, we introduce constraints that take the different room capacities and demands of the courses into account. Let set \mathcal{S} contain the different room capacities. Let $\mathcal{C}_{\geq s}$ denote all courses with demand larger than s ; and $\mathcal{R}_{\geq s}$ denotes rooms with capacity more than s seats. For each $s \in \mathcal{S}$, except the biggest, and for all $c \in \mathcal{C}_{\geq s}$ there is a three-indexed binary variable $y_{s,c,p}$. If all rooms would dispose of different capacities, we would introduce here another three-indexed problem formulation based on a course/time/room assignment. But if many rooms share the same capacity, which is often the case in real-world problems, the amount of the three-indexed $y_{s,c,p}$ variables is reasonably large. So for real-world problems this constraint can be integrated in our decomposition model without breaking the basic idea of this approach. We then add for all $p \in \mathcal{P}$ and all $s \in \mathcal{S}$

$$x_{c,p} - y_{s,c,p} \geq 0 \quad c \in \mathcal{C}_{\geq s}, \quad (7)$$

$$\sum_{c \in \mathcal{C}_{\geq s}} x_{c,p} - y_{s,c,p} \leq |\mathcal{R}_{\geq s}|. \quad (8)$$

Variable $y_{s,c,p}$ is set to one if course c takes place in a room of capacity smaller than s . By constraint (8) we ensure that this does not happen for more courses than we have rooms of appropriate capacity; otherwise, we incur a penalty which is considered in the objective

function. Variable $y_{s_i,c,p}$ receives the coefficient $obj_{s_i,c,p}$ which mainly reflects the difference between the demand $dem(c)$ of course c and s_i . We define $obj_{s_i,c,p}$ as:

$$obj_{s_i,c,p} := \min\{dem(c) - s_i, s_{i+1} - s_i\}. \tag{9}$$

We add to the objective function (1)

$$\sum_{c \in \mathcal{C}_{\geq s}} obj_{s,c,p} \cdot y_{s,c,p}. \tag{10}$$

3.2 MinimumWorkingDay constraints

For each course c we specify a minimum number $mnd(c)$ of days, among which its lectures should be distributed. This constraint goes into the first decomposition stage. We introduce a binary variable $z_{c,d}$ for every course c and every eligible day d for this course. Now we add

$$\sum_{p \in d} x_{c,p} - z_{c,d} \geq 0 \quad \forall c \in \mathcal{C}, d \in \mathcal{D}. \tag{11}$$

So, $z_{c,d}$ can be set to one only if course c takes place at some period of day d . Furthermore, we introduce another integer variable w_c and the following constraint:

$$\sum_{d \in \mathcal{D}} z_{c,d} + w_c \geq mnd(c) \quad \forall c \in \mathcal{C}. \tag{12}$$

Obviously, variable w_c may take value zero only if course c takes place on more than $mnd(c) - 1$ days. According to the penalty system introduced in Di Gaspero et al. (2007), McCollum et al. (2009) we add to the objective function (1)

$$\sum_{c \in \mathcal{C}} 5 \cdot w_c. \tag{13}$$

3.3 CurriculumCompactness constraints

For every curriculum, the corresponding courses should take place consecutively over a day. We will see that, even though easily incorporated in the first stage, these soft constraints have a negative influence on solution times. For every period $p \in \mathcal{P}$ and every curriculum $cu \in \mathcal{CU}$ we introduce a binary variable $r_{p,cu}$ and the following constraint:

$$\sum_{c \in cu} x_{c,p} - r_{cu,p} = 0 \quad \forall cu \in \mathcal{CU}, p \in \mathcal{P}. \tag{14}$$

Variable $r_{cu,p}$ assumes value one if some course of curriculum cu takes place at period p , and zero otherwise. Note that constraints (14) imply the stable set conditions (4) for curriculum conflicts. Again with the help of binary indicator variables $v_{cu,p}$ we model curriculum compactness:

$$-r_{cu,p-1} + r_{cu,p} - r_{cu,p+1} - v_{cu,p} \leq 0 \quad \forall cu \in \mathcal{CU}, p \in \mathcal{P}. \tag{15}$$

If period p is the last of the day the term $r_{cu,p+1}$ is omitted, and if p is the first period of the day the term $r_{cu,p-1}$ is omitted. Obviously, $v_{cu,p}$ has to be set to one if the curriculum

cu has an isolated lecture at period p . Consequently, the following term is added to the objective (1):

$$\sum_{cu \in \mathcal{CU}, p \in \mathcal{P}} 2 \cdot v_{cu,p}. \quad (16)$$

3.4 RoomStability constraints

Room stability encourages all lectures of a course to take place in the same room. In contrast to the previous soft constraints, we currently see no way to respect this already in the first stage. As a consequence, the perfect matching structure of the second stage is destroyed, in particular integrality of solutions is lost, and we have to resort to integer programming. The negative impact on running times is significant.

As will be seen in Sect. 3.6 the IP Formulation of the second stage still resembles the standard matching formulation on bipartite graphs. We introduce binary variables $u_{c,p}v_{r,p}$ which assume value one iff course c takes place in room r at period p . Furthermore, we add binary variables $y_{c,r}$ for each course c and each eligible room r , which are included via

$$\sum_{p \in \mathcal{P}} u_{c,p}v_{r,p} - |\mathcal{P}| \cdot y_{c,r} \leq 0 \quad \forall c \in \mathcal{C}, \forall r \in \mathcal{R}. \quad (17)$$

Variable $y_{c,r}$ must assume value one, if course c takes place in room r at least once. The second stage objective function reads

$$\sum_{c \in \mathcal{C}, r \in \mathcal{R}} y_{c,r}. \quad (18)$$

Clearly, if (18) is minimized over all feasible course/room assignments, the RS constraint is fulfilled best possibly according to the underlying bipartite graph. But as we will see, the bipartite graph depends on the solution of the first decomposition stage. It is therefore possible (and it happens) that the obtained solution is not a globally optimal one.

3.5 IP Formulation for the first stage

The introduction of soft constraints resulted in a significantly altered model as compared to (1)–(5), not only visibly but also in terms of combinatorial structures. It turns out that this has a negative impact on computation times. The only constraint we did not yet take care of is that no two courses by the same teacher may be scheduled in parallel. Denote by \mathcal{T} the set of teachers, and by $C(t)$ the courses given by teacher $t \in \mathcal{T}$.

$$\begin{aligned} \min \quad & \sum_{p \in \mathcal{P}, s \in \mathcal{S}, c \in \mathcal{C}_{\geq s}} \text{obj}_{s,c,p} \cdot y_{s,c,p} + \sum_{c \in \mathcal{C}} 5 \cdot w_c + \sum_{cu \in \mathcal{CU}, p \in \mathcal{P}} 2 \cdot r_{cu,p} \\ \text{subject to} \quad & \sum_{p \in \mathcal{P}} x_{c,p} = l(c) \quad \forall c \in \mathcal{C}, \\ & \sum_{c \in \mathcal{C}} x_{c,p} \leq |\mathcal{R}| \quad \forall p \in \mathcal{P}, \\ & x_{c,p} - y_{s,c,p} \geq 0 \quad \forall s \in \mathcal{S}, c \in \mathcal{C}_{\geq s}, p \in \mathcal{P}, \\ & \sum_{c \in \mathcal{C}_{\geq s}} x_{c,p} - y_{s,c,p} \leq |\mathcal{R}_{\geq s}| \quad \forall s \in \mathcal{S}, p \in \mathcal{P}, \\ & \sum_{p \in \mathcal{D}} x_{c,p} - z_{c,d} \geq 0 \quad \forall c \in \mathcal{C}, d \in \mathcal{D}, \end{aligned}$$

$$\begin{aligned}
 \sum_{d \in \mathcal{D}} z_{c,d} + w_c &\geq mnd(c) && \forall c \in \mathcal{C}, \\
 \sum_{c \in \mathcal{CU}} x_{c,p} - r_{cu,p} &= 0 && \forall cu \in \mathcal{CU}, p \in \mathcal{P}, \\
 -r_{cu,p-1} + r_{cu,p} - r_{cu,p+1} - v_{cu,p} &\leq 0 && \forall cu \in \mathcal{CU}, p \in \mathcal{P}, \\
 \sum_{c \in \mathcal{C}(t)} x_{c,p} &\leq 1 && \forall t \in \mathcal{T}, p \in \mathcal{P}, \\
 r_{cu,p} &\in \{0, 1\}, \\
 v_{cu,p} &\in \{0, 1\}, \\
 w_c &\in \mathbb{Z}_+, \\
 x_{c,p} &\in \{0, 1\}, \\
 y_{s,c,p} &\in \{0, 1\}, \\
 z_{c,d} &\in \{0, 1\}.
 \end{aligned}$$

3.6 IP formulation for the second stage

Originally, the second stage was to solve a minimum cost perfect matching problem for each period. The situation is more involved in light of the soft constraints. Let $G = (U \cup V, E)$ be a bipartite graph with node set $U \cup V$ defined according to the values $x_{c,p}^*$ of variables $x_{c,p}$ obtained in the first stage. Let $cap(r)$ denote the capacity of room r and $dem(c)$ denote the seat demand of course c . Given a solution x^* the graph G is defined as follows for each $p \in \mathcal{P}$:

$$U = \{u_{c,p} : x_{c,p}^* = 1, c \in \mathcal{C}\},$$

$$V = \{v_{r,p} : r \in \mathcal{R}\},$$

$$E = \begin{cases} u_{c,p}v_{r,p} & \text{if } y_{s,c,p} = 0 \text{ and } dem(c) \leq cap(r), \\ u_{c,p}v_{r,p} & \text{if } y_{s,c,p} = 1, dem(c) > cap(r), cap(r) = \max\{cap(\hat{r}) : cap(\hat{r}) < dem(c)\}. \end{cases}$$

We denote for $x \in U \cup V$ by $\delta(x) = \{e \in E : \exists y \in U \cup V, e = xy \vee e = yx\}$ the cut of x in G . Then, the integer program for the second stage reads as

$$\begin{aligned}
 \min \quad & \sum_{c \in \mathcal{C}, r \in \mathcal{R}} y_{c,r} \\
 \text{subject to} \quad & \sum_{p \in \mathcal{P}} u_{c,p}v_{r,p} - |\mathcal{P}| \cdot y_{c,r} \leq 0 \quad \forall c \in \mathcal{C}, r \in \mathcal{R}, \tag{19}
 \end{aligned}$$

$$\sum_{u_{c,p}v_{r,p} \in \delta(u_{c,p})} u_{c,p}v_{r,p} = 1 \quad \forall u_{c,p} \in U, \tag{20}$$

$$\sum_{u_{c,p}v_{r,p} \in \delta(v_{r,p})} u_{c,p}v_{r,p} \leq 1 \quad \forall v_{r,p} \in V, \tag{21}$$

$$u_{c,p}v_{r,p} \in \{0, 1\},$$

$$y_{r,p} \in \{0, 1\}.$$

The constraints consist of two different parts. The RS constraints are given in (19), cf. (17). Constraints (20) and (21) are from the standard formulation of a (one-sided perfect) matching in a bipartite graph, ensuring that each course gets one room assigned in a

period when it takes place (20), and that no room is occupied more than once at the same time (21).

4 Extensions

In Di Gaspero et al. (2007) and McCollum et al. (2009) several more constraints are mentioned which are relevant in practice, but do not appear in the ITC2007 competition's problem definition for the purpose of a cleaner presentation. The authors state that "if in the future this formulation will prove to be inappropriate (e.g., too simple), some features could be reintroduced for future research." In this section we demonstrate how to incorporate all of them into our model; some experience is given in Sect. 5.

It is an advantage of our decomposition approach that several constraints, in particular those relating to rooms, are easily dealt with, some are even automatically satisfied. Conditions depending on the curriculum can be modeled via the $r_{cu,p}$ variables but require new constraints in the decomposition's first stage IP formulation from Sect. 3.5.

4.1 Lunch break for students

For each curriculum cu and a day d let p_1, p_2 be the periods around noon. Then we add the following constraint:

$$r_{cu,p_1} + r_{cu,p_2} - l_{cu,d} \leq 1. \quad (22)$$

If curriculum cu cannot have a lunch break, because courses are scheduled around noon on day d , the binary variable $l_{cu,d}$ has to be set to one. This is penalized in the objective function with two units per violation.

4.2 Specific patterns in curriculum compactness

This soft constraint is only sloppily defined in Di Gaspero et al. (2007) and McCollum et al. (2009), but individually penalizing specific patterns of non-contiguous lectures of courses in a curriculum can be done by encoding them similarly to the pattern in constraint (15).

4.3 Curriculum dependent maximum student dayload

The maximal number $dload$ of courses a student should take in a given curriculum cu per day d can be softly limited in the same way as we encourage lunch breaks. Let p_1, \dots, p_k be the periods of day d , then we add a constraint

$$\sum_{i=1}^k r_{cu,p_i} - dl_{cu,d} \leq dload. \quad (23)$$

The integer variable $dl_{cu,d}$ assumes a strictly positive value if the maximum dayload is exceeded. Every violation is penalized with four units.

4.4 Consecutiveness of lectures

Some lectures have to be (or must not be) scheduled in consecutive periods. Two parts of the formulation need to be changed. The stable set conditions (4) based on the conflict graph can be adapted straight forwardly. It is more complicated, yet doable, to adjust Hall's conditions (3), but the discussion is too involved for the scope of this paper.

4.5 Room unavailability

If a room is not available at some period p , this room simply does not appear in the corresponding bipartite graph G_p , and is omitted in the Hall's conditions (3) or equivalent constraints for this period.

4.6 Appropriate room sizes

A lecture should not take place in a too large room. This requirement is symmetric to the room capacity constraints, and is modeled in an analogous way. Again, let \mathcal{S} be the set of different room capacities. For all except the largest $s \in \mathcal{S}$ we introduce further constraints. By $\mathcal{C}_{\leq s}$ we denote all courses with demand no greater than s , and by $\mathcal{R}_{\leq s}$ denote the rooms with capacity no greater than s . Given $s \in \mathcal{S}$, for all $c \in \mathcal{C}_{\leq s}$ we introduce a binary variable $t_{s,c,p}$ with meaning symmetric to variables $y_{s,c,p}$ in Sect. 3.1. We add constraints

$$x_{c,p} - t_{s,c,p} \geq 0 \quad \forall s \in \mathcal{S} \quad c \in \mathcal{C}_{\leq s} \quad (24)$$

$$\sum_{c \in \mathcal{C}_{\leq s}} x_{c,p} - t_{s,c,p} \leq |\mathcal{R}_{\leq s}| \quad \forall s \in \mathcal{S} \quad (25)$$

A penalty reflecting the difference between s and the seat demand of course c is incurred for using $t_{s,c,p}$.

4.7 Complex weights for soft constraint violations

By our use of binary indicator variables for each individual violation of a soft constraint (that is, for each single curriculum, day, period, or room) we may give individual penalties, in particular depending on the number of students which take a given course.

4.8 Teacher preferences

Teachers may express priorities reflecting when they prefer (not) to teach. This is the original objective function used e.g., at TU Berlin; we formulated this objective in Sect. 2.

5 Computational study

We report on three different sets of experiments. In the first (Sect. 5.1), we deal with “the Udine instances,” in particular those used at ITC2007. The second set (Sect. 5.2) contains somewhat larger instances from a recent paper by Cesco et al. (2008). For both sets we consider both, the “basic” formulation (Di Gaspero and Schaerf 2006) (without RS constraints), and the “extended” formulation (Di Gaspero et al. 2007; McCollum et al. 2009) with all four types of soft constraints. The final (smallest) set (Sect. 5.3) contains much larger instances with only hard constraints. This last set reflects the timetabling situation at the Technical University of Berlin. All experiments were run on a 3.4 GHz Linux PC with 1 GB memory; unless specified otherwise, we solved integer programs with CPLEX 11.0.1. The reported optimality gaps were computed relative to the upper bound, i.e., as (upper bound–lower bound) / upper bound.

The curriculum-based course timetabling web site <http://tabu.diegm.uniud.it/ctt/> is most helpful in making results comparable. First of all, they offer a solution validator which we

used, of course, to validate our results (solution files can be requested from the authors by email). From the same web site one can download a program to benchmark machine speed. In our computations, one *CPU time unit* corresponds to the time allowed for one run in the ITC2007 competition: This should be around 400 seconds on a reasonable PC. For ITC2007, the program of every finalist was run 10 times, each time with a different random seed. Thus, it took 10 CPU time units to achieve their respective best solutions. Sometimes, the competition winner Tomáš Müller (2008) did not achieve the overall best result for an instance. Since we also compare ourselves against these overall best solutions of all of the five finalists, we say that it took 5·10 CPU time units to obtain these solutions. When we compare ourselves to the best solutions by the university of Udine's Scheduling and Timetabling Group (SaTT) we assumed they used 40 CPU time units since they started 40 runs to obtain their best results. Since in contrast our approach is entirely deterministic, it is fair to allow ourselves a solution time equivalent to what is used in total in the respective runs of these various groups.

There are several (similar) tables, and if you are in a rush, the most important conclusions can be drawn from Tables 2, 4, 6, and 8.

5.1 The Udine benchmark instances

5.1.1 The original benchmarks from Di Gaspero and Schaerf (2003), Di Gaspero and Schaerf (2006)

In Table 1 we list for the first time proven optimal solutions for *all the four* instances used in Di Gaspero and Schaerf (2003, 2006), in particular *test4* was unsolved within the given time limits. These original instances do not feature RS constraints.

For all except the last instance, running times are quite short. Taking into account that no previous approach has produced optimal results for all four instances, this is remarkable and demonstrates the usefulness of our approach. Among all soft constraints, curriculum compactness (CC) appears to destroy the combinatorial structure of the timetabling problem the most. An impressive proof for this is given in Table 1 where these constraints are dropped.

5.1.2 The role of the solver

It should be mentioned that the last years have seen great improvements in integer programming solvers, so one might suspect that our ability to solve *test1–4* is mainly due to this fact; however, with the several years old CPLEX9 we are able to produce optimal solutions

Table 1 Optimal solutions values for the Udine problem instances (basic formulation Di Gaspero and Schaerf 2006). We list instance names, our objective function values (soft constraint penalties), and the CPU time needed for computations. On the right the we see the results when the CC constraints are omitted

Instance	Basic formulation (Di Gaspero and Schaerf 2006)		Without CC constraint	
	obj	CPU sec.	obj	CPU sec.
<i>test1</i>	212	15.40	200	0.14
<i>test2</i>	8	6.31	0	0.08
<i>test3</i>	35	82.33	5	0.11
<i>test4</i>	27	1607.30	0	0.17

to the first three instances within computation times comparable to those in Table 1, and a very good solution for `test4` (value 29) in about an hour. However, actually proving this quality is not possible with CPLEX9, since the lower bound does not improve at all (the zero-half cuts of later CPLEXes do help a lot in this respect).

In order to check the necessity of a commercial solver in the first place we tested the non commercial, open source solvers SCIP (Achterberg 2009) (scip.zib.de) and CBC (www.coin-or.org/Cbc/) to solve our integer programs. These could not match the good running times of the commercial solver CPLEX. The use of a commercial solver (and thus, the possible lack of reproducibility of results on any machine) is, in fact, the reason why we did not submit our results to the ITC2007 competition.

5.1.3 Benchmarks from ITC2007

The second International Timetabling Competition, ITC2007, extended the definition of soft constraints by adding room stability (RS). Seven instances (`comp01–07`) were provided at the outset of the competition, seven more (`comp08–14`) followed closer to the deadline (and seven more after the deadline, but these are not yet available to us). Table 2 lists our results. As one can see we are always (except twice) better than the average run of the ITC2007 winner, and we are very competitive with the respective best results obtained by all the five finalists. Results obtained in Table 2 are with CPLEX' zero-half cuts turned on.

Table 2 We compare ourselves against the university of Udine's Scheduling and Timetabling Group (SaTT), against the objective of the ITC2007 winner, averaged over all his 10 runs, against his respective best run, and against the overall best run of all the five finalists

Instance	Basic formulation (Di Gaspero and Schaefer 2006)		Extended formulation (Di Gaspero et al. 2007; McCollum et al. 2009)			SaTT	us
	SaTT	us	ITC2007				
			winner avg	winner best	finalists best		
<code>comp01</code>	4	4	5.0	5	5	5	13
<code>comp02</code>	35	31	61.3	51	50	75	43
<code>comp03</code>	52	42	94.8	84	71	93	76
<code>comp04</code>	21	18	42.8	37	35	45	38
<code>comp05</code>	244	253	343.5	330	309	326	314
<code>comp06</code>	27	16	56.8	48	48	62	41
<code>comp07</code>	13	3	33.9	20	20	38	19
<code>comp08</code>	24	20	46.5	41	40	50	43
<code>comp09</code>	61	59	113.1	109	105	119	102
<code>comp10</code>	10	8	21.3	16	16	27	14
<code>comp11</code>	0	0	0.0	0	0	0	0
<code>comp12</code>	268	316	351.6	333	333	358	405
<code>comp13</code>	38	33	73.9	66	66	77	68
<code>comp14</code>	30	29	61.8	59	57	59	54
CPU time units	40	10	1	10	50	40	10

Table 3 Computation times for ITC2007 (extended formulation Di Gaspero et al. 2007; McCollum et al. 2009) listed separately for the two decomposition stages (in seconds), for time limits of a total of 1, 10, and 40 CPU time units in the different subtables (a)–(c) (detailed limits on the two stages are given in the respective headings). We report (in that order) the instance name, the objective function value (plus lower bound and optimality gap), and the time to reach that solution. As a reference we also give the objective value of and the time to reach the first feasible integer solution. Information on the second stage is similar. Computation times listed in this table are rather coarsely reported and only serve as an indicator. CPLEX11 is used with default parameter settings

(a) Overall time limit: 1 CPU time unit

Instance	stage 1 (time limit: 300 sec.)						stage 2 (time limit: 80 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
comp01	4	4.00	0.00	<5	4	<5	8	<1	8	<1	12
comp02	273	0.00	100.00	120	430	<50	16	20	147	2	239
comp03	191	0.00	100.00	261	468	<10	3	60	143	10	194
comp04	36	21.90	39.15	264	358	<5	8	40	144	10	44
comp05	956	91.83	90.39	290	1241	<90	9	<1	16	<1	965
comp06	346	7.00	97.98	280	541	<100	49	80	180	3	395
comp07	448	0.00	100.00	290	525	<190	68	80	225	3	525
comp08	39	29.20	25.11	190	344	<4	39	70	173	4	78
comp09	113	36.89	67.35	290	444	<2	2	80	160	2	115
comp10	194	2.00	98.97	200	425	<110	41	60	207	2	235
comp11	0	0.00	0.00	<1	0	<1	7	<1	7	<1	7
comp12	1119	28.08	97.49	290	1119	290	3	4	77	<1	1122
comp13	75	32.17	57.11	270	492	<3	23	80	161	2	98
comp14	110	39.50	71.79	290	449	<20	3	80	141	1	113

(b) Overall time limit: 10 CPU time units

Instance	stage 1 (time limit: 3300 sec.)						stage 2 (time limit: 500 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
comp01	4	4.00	0.00	<5	4	<5	8	<1	8	<1	12
comp02	93	8.00	91.40	~3000	430	<120	0	208	138	2	93
comp03	84	0.00	100.00	3140	468	<40	2	300	132	10	86
comp04	35	27.43	21.61	2960	358	<5	5	330	145	10	41
comp05	463	24.30	95.26	2800	1241	<430	5	<1	13	<1	468
comp06	66	10.00	84.85	~3000	541	<300	13	300	181	3	79
comp07	8	2.00	75.00	~2000	525	<360	20	413	234	3	28
comp08	37	34.00	8.11	2990	344	<10	11	200	177	4	48
comp09	106	41.00	60.79	3280	444	<2	0	439	169	2	106
comp10	4	4.00	0.00	2385	425	<220	40	130	207	2	44
comp11	0	0.00	0.00	<1	0	<1	7	<1	7	<1	7
comp12	657	31.28	95.24	~2500	1119	290	0	4	81	<1	657
comp13	61	38.60	36.72	~1930	492	<3	6	300	155	3	67
comp14	51	41.00	18.66	~1500	449	<20	3	284	146	1	54

Table 3 (Continued)

(c) Overall time limit: 40 CPU time units

Instance	stage 1 (time limit: 13000 sec.)						stage 2 (time limit: 2200 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
comp01	4	4.00	0.00	<5	4	<5	8	<1	8	<1	12
comp02	45	10.33	77.04	~11500	430	<120	1	2000	151	2	46
comp03	66	25.00	100.00	~12500	468	<40	0	432	127	10	66
comp04	35	27.43	21.61	2960	358	<5	3	1300	51	10	38
comp05	365	107.97	70.43	12700	1241	<430	3	<1	38	<1	368
comp06	37	10.00	72.97	7526	541	<300	14	2000	187	3	51
comp07	6	6.00	0.00	10000	525	<360	19	2000	221	3	25
comp08	37	37.00	0.00	500	344	<10	8	1200	178	4	44
comp09	99	45.89	53.65	12800	444	<2	0	500	165	2	99
comp10	4	4.00	0.00	2385	425	<220	12	2000	207	2	16
comp11	0	0.00	0.00	<1	0	<1	7	<1	7	<1	7
comp12	546	52.70	90.34	11000	1119	290	1	4	79	<1	548
comp13	61	40.81	33.72	~1930	492	<3	5	800	155	3	66
comp14	51	45.94	9.92	~1500	449	<20	2	900	146	1	53

In Table 3 we list statistics separately for the two stages of the decomposition for various overall time bounds. These results were obtained with CPLEX11 default parameter settings.

5.1.4 Lower bounds

Meta heuristics are powerful to produce solutions to quite large timetabling instances. However, assessing the quality of these solutions is much harder. Recently, Burke et al. (2008a) proposed a branch-and-cut algorithm to obtain lower bounds for the ITC 2007 instances. We note that the time to solve our linear programming (LP) relaxation is much smaller since our formulation contains much fewer constraints and variables (about a factor of two to three in each dimension). Further, the program presented in Burke et al. (2008a) is not yet suited to produce feasible integer solutions; this is why Burke et al. resorted to heuristics for this. We list the lower bounds obtained by our approach in Table 4.

5.1.5 Extensions

In Sect. 4 we discussed several extensions for soft constraints as proposed in Di Gaspero et al. (2007) and McCollum et al. (2009). Table 5 lists our results for the original Udine and the first seven ITC2007 instances, when the problem definition is exemplarily extended by the *Maximum Dayload* and the *Lunch Break* constraints. We did not include the other extended soft constraints in this study.

5.2 Instances with more courses

A hint on the potential of our approach when applied to larger instances is given on data recently introduced by Cesco et al. (2008). Some of them have a (slightly) larger number of courses (DDS1 and DDS4), and our integer program performs significantly better than what was previously known, see Table 6.

Table 4 Lower bounds (LBs) obtained by Burke et al. (2008a) and with our approach (first stage) for the extended formulation (Di Gaspero et al. 2007; McCollum et al. 2009). On the left one can see the computation times to solve the LP relaxation, then the LBs in the root node with our plain formulation, and after adding CPLEX' zero-half cuts; finally, the LBs after half an hour computation time. Numbers are taken from a first draft of (Burke et al. 2008a); updated results are not available to us

Instance	Root relaxation (sec.)		LB (root)		LB (after 30 min.)	
	(Burke et al. 2008a)	us	us	us w/cuts	(Burke et al. 2008a)	us
comp01	3.58	0.09	4	4	5	4
comp02	54.90	0.68	0	0	6	8
comp03	49.97	0.59	0	1	43	23
comp04	41.00	0.12	0	11.5	2	26.27
comp05	84.64	1.97	17	92.45	183	100.9
comp06	73.17	0.88	6	7	6	7
comp07	192.35	1.47	0	0	0	0
comp08	43.25	0.61	0	1.16	2	33.2
comp09	48.23	0.51	0	18.2	0	39.84
comp10	105.03	1.00	0	2	0	3.91
comp11	7.69	0.13	0	0	0	0
comp12	134.76	2.92	3	30.25	5	31.29
comp13	37.34	0.67	0	20	0	37
comp14	55.86	0.72	0	39.5	0	41

Table 5 Best solutions for the instances from Di Gaspero and Schaerf (2003, 2006) and the first seven from ITC2007, with extensions as discussed in Sects. 4 and 5.1.5. Bold face marks optimal solutions

Instance	Obj	Lower bound	Gap	Status	CPU sec.
test1	217	215	0.97%	feasible	150
test2	59	59	0.00%	optimal	26.23
test3	127	127	0.00%	optimal	125
test4	48	45.47	5.25%	feasible	3600
comp01	8	8	0.00%	optimal	11.42
comp02	417	35.71	92.12%	feasible	3600
comp03	202	59	70.07%	feasible	3600
comp04	28	28	0.00%	optimal	1183
comp05	418	120.73	71.12%	feasible	3600
comp06	96	11.08	88.45%	feasible	3600
comp07	407	3	99.26%	feasible	3600

5.3 Simulated data from Technical University Berlin

As we have said, our original motivation was to keep hard constraints, if this is possible. At the Technical University of Berlin, room capacities are considered hard, and a number of features have to be provided by a room if requested by a course (Internet access, PC/projector, blackboard, location, etc.). This gives a much larger number of different room types, but in general fewer eligible rooms per course. All other soft constraints presented here are not

Table 6 Slightly larger instances taken from De Cesco et al. (2008); our approach compared to the university of Udine's Scheduling and Timetabling Group (SaTT); bold face indicates optimal solutions. Zero-half cuts are turned on in these computations

Instance	Basic formulation (Di Gaspero and Schaerf 2006)		Extended formulation (Di Gaspero et al. 2007; McCollum et al. 2009)	
	SaTT	us	SaTT	us
	DDS1	238	39	1024
DDS2	0	0	0	0
DDS3	0	0	0	0
DDS4	233	19	261	68
DDS5	0	0	0	0
DDS6	5	0	11	4
DDS7	0	0	0	0
CPU time units	40	10	40	10

respected, as they are not relevant for this university. Since the used timetabling database is in an incomplete and inconsistent state, we decided to develop a simulation tool which is able to create large problem instances with near real-world character. The instances are available at www.math.tu-berlin.de/lach/.

We present statistics of three representative instances of different sizes, cf. Table 8. The key data (not listed here) of the large instance is almost identical to that of Technical University of Berlin (which is a rather large university). We give running times for a preprocessing step necessary to generate only the actually needed Hall conditions (3), and for the two decomposition stages. These times are acceptable, even though for an interactive timetable design, some tuning would be necessary.

6 Perspectives

Formulating a problem as an integer program (IP) is a very powerful modeling tool for formal and precise identification of relevant decisions and constraints. In our view, this has been the primary use of integer programming in university course timetabling so far, rather than actually trying to compute optimal solutions with its help. Only recently, researchers started to exploit the problem's structure, as in we did in this paper, in order to design non-standard algorithms for solving the resulting IPs. A recent example besides our decomposition approach presented here is the branch-and-cut algorithm for obtaining lower bounds in Burke et al. (2008a). Actually, we believe that optimal or high quality solutions for university course timetabling, also with the enormous progress made in IP solver technologies, will not be computed with off-the-shelf solvers, but will call for tailored algorithms. This requires to further investigate the combinatorial structure hidden in the problem, e.g., in the soft constraints, and we are encouraged by our good results to pursue such studies in the future.

Our approach was originally meant to solve instances from Berlin's technical university where all constraints are considered hard, see Lach and Lübbecke (2008). However, feedback on that approach motivated us to evaluate its suitability for incorporating soft constraints, or to obtain lower bounds. We believe it is fair to say that our algorithm certainly is not best in all possible situations, however, it competes quite well for several different

Table 7 Solution statistics reported separately for the two decomposition stages for the DDS1–7 instances, in the same way as we did in Table 3. Default parameter settings are used

(a) Overall time limit: 1 CPU time unit

Instance	stage 1 (time limit: 300 sec.)						stage 2 (time limit: 80 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
DDS1	147	42.99	70.75	<240	237	<200	144	60	499	<5	291
DDS2	0	0.00	0.00	<1	60	<1	0	<5	47	<5	0
DDS3	0	0.00	0.00	<1	22	<1	0	60	78	<5	0
DDS4	675	15.00	97.78	300	1067	<300	376	40	484	<5	1051
DDS5	0	0.00	0.00	22	147	<2	42	70	298	<5	42
DDS6	159	0.00	0.00	280	460	<100	27	80	171	<5	186
DDS7	0	0.00	0.00	52	127	<3	4	80	110	<5	4

(b) Overall time limit: 10 CPU time units

Instance	stage 1 (time limit: 3300 sec.)						stage 2 (time limit: 500 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
DDS1	48	48.00	0.00	2000	237	<200	55	800	499	<5	103
DDS2	0	0.00	0.00	<1	60	<1	0	<5	47	<5	0
DDS3	0	0.00	0.00	<1	22	<1	0	60	78	<5	0
DDS4	17	15.00	11.76	700	1067	<300	95	500	484	<5	112
DDS5	0	0.00	0.00	22	147	<2	10	800	298	<5	10
DDS6	4	0.00	100.00	1000	460	<100	5	500	171	<5	9
DDS7	0	0.00	0.00	52	127	<3	0	309	110	<5	0

(c) Overall time limit: 40 CPU time units

Instance	stage 1 (time limit: 13000 sec.)						stage 2 (time limit: 2200 sec.)				total obj
	obj	LB	gap%	time	obj 1st	time 1st	obj	time	obj 1st	time 1st	
DDS1	48	48.00	0.00	2000	237	<200	35	1800	499	<5	83
DDS2	0	0.00	0.00	<1	60	<1	0	<5	47	<5	0
DDS3	0	0.00	0.00	<1	22	<1	0	60	78	<5	0
DDS4	17	15.00	11.76	700	106	<300	75	2000	484	<5	92
DDS5	0	0.00	0.00	22	147	<2	10	700	298	<5	10
DDS6	0	0.00	0.00	3000	460	<100	3	2000	171	<5	3
DDS7	0	0.00	0.00	52	127	<3	0	309	110	<5	0

purposes, as is demonstrated by our computational study. This is even more true since we do not use any particular tuning to the instances or situation (lower or upper bound). Additionally, we are capable of computing feasible solutions which respect constraints which are listed in the “research agenda in automated timetabling” (McCollum et al. 2009) already today. In that sense, the proposed approach may serve as a very *robust* starting point for more ambitious goals in timetabling.

From a practical point of view, one is interested in warm-starting computations from previous timetables in such a way, that small changes in the input result in small changes

Table 8 Statistics and computation times for the simulated instances according to Technical University of Berlin's course database

Instance	Courses	Lectures	Rooms	Violations	Preproc.	Stage 1	Stage 2
small	180	420	35	0	45 sec.	9 sec.	3 sec.
medium	950	2100	165	0	307 sec.	52 sec.	6 sec.
large	2100	4640	345	0	1235 sec.	5106 sec.	5 sec.

in the constructed timetable. This other kind of *robustness* could be considered already in constructing the first timetable via the framework of robust optimization; however, this will require entirely new research efforts and is beyond the scope of this paper.

References

- Achterberg, T. (2009). SCIP: solving constraint integer programs. *Mathematical Programming Computation*, 1, 1–41.
- Burke, E. K., & Petrovic, S. (2002). Recent research directions in automated timetabling. *European Journal of Operational Research*, 140(2), 266–280.
- Burke, E. K., Mareček, J., Parkes, A. J., & Rudová, H. (2007). *On a clique-based integer programming formulation of vertex colouring with applications in course timetabling*. Technical Report NOTTCS-TR-2007-10, The University of Nottingham. [arXiv:0710.3603v2](https://arxiv.org/abs/0710.3603v2).
- Burke, E. K., Mareček, J., Parkes, A. J., & Rudová, H. (2008a). A branch-and-cut procedure for the Udine course timetabling problem. In E.K. Burke & M. Gendreau (Eds.), *Proceedings of the 7th international conference on the practice and theory of automated timetabling, PATAT 2008*, Montréal, CA.
- Burke, E. K., Mareček, J., Parkes, A. J., & Rudová, H. (2008b). Penalising patterns in timetables: Strengthened integer programming formulations. In J. Kalcsics & S. Nickel (Eds.) *Operations research proceedings 2007* (pp. 409–414). Berlin: Springer.
- Burke, E. K., Mareček, J., Parkes, A. J., & Rudová, H. (2008). *Uses and abuses of MIP in course timetabling*. Poster at the workshop on mixed integer programming, MIP2007, Montréal, 2008. Available online at <http://cs.nott.ac.uk/jxm/timetabling/mip2007-poster.pdf>.
- Carter, M. W. (2001). A comprehensive course timetabling and student scheduling system at the University of Waterloo. In E. Burke & W. Erben (Eds.) *Lect. Notes Comp. Science: Vol. 2079. Proceedings of the 3th international conference on the practice and theory of automated timetabling, PATAT 2000* (pp. 64–82). Berlin: Springer.
- Daskalaki, S., & Birbas, T. (2005). Efficient solutions for a university timetabling problem through integer programming. *European Journal of Operational Research*, 127(1), 106–120.
- Daskalaki, S., Birbas, T., & Housos, E. (2004). An integer programming formulation for a case study in university timetabling. *European Journal of Operational Research*, 153, 117–135.
- De Cesco, F., Di Gaspero, L., & Schaerf, A. (2008). Benchmarking curriculum-based course timetabling: Formulations, data formats, instances, validation, and results. In E.K. Burke & M. Gendreau (Eds.), *Proceedings of the 7th international conference on the practice and theory of automated timetabling, PATAT 2008*, Montréal, CA.
- Di Gaspero, L., & Schaerf, A. (2003). Multi neighborhood local search with application to the course timetabling problem. In E. Burke & P. De Causmaecker (Eds.) *Lect. Notes Comp. Science: Vol. 2740. Proceedings of the 4th international conference on the practice and theory of automated timetabling, PATAT 2002* (pp. 262–275). Berlin: Springer.
- Di Gaspero, L., & Schaerf, A. (2006). Neighborhood portfolio approach for local search applied to timetabling problems. *Journal of Mathematical Modelling and Algorithms*, 5, 65–89.
- Di Gaspero, L., McCollum, B., & Schaerf, A. (2007). *The second international timetabling competition (ITC-2007): Curriculum-based course timetabling (track 3)*. Technical Report QUB/IEEE/Tech/ITC2007/CurriculumCTT/v1.0, Queen's University, Belfast, United Kingdom, August 2007.
- Lach, G., & Lübbecke, M. E. (2008). Optimal university course timetables and the partial transversal polytope. In C. C. McGeoch (Ed.) *Lect. Notes Comput. Sci.: Vol. 5038. Proceedings of the 7th workshop on experimental algorithms (WEA)* (pp. 235–248). Berlin: Springer.

- Lovász, L., & Plummer, M. D. (1986). *Matching theory*. Amsterdam: North-Holland.
- McCollum, B., McMullan, P., Paechter, B., Lewis, R., Schaerf, A., Di Gaspero, L., Parkes, A. J., Qu, R., & Burke, E. (2009, in press). Setting the research agenda in automated timetabling: The second international timetabling competition. *INFORMS Journal on Computing*.
- Müller, T. (2008). ITC2007 solver description: A hybrid approach. In *Proceedings of the 7th international conference on the practice and theory of automated timetabling, PATAT 2008*, Montréal, CA.
- Qualizza, A., & Serafini, P. (2005). A column generation scheme for faculty timetabling. In E. K. Burke, & M. A. Trick (Eds.) *Lect. Notes Comp. Science: Vol. 3616. Proceedings of the 5th international conference on the practice and theory of automated timetabling, PATAT 2004* (pp. 161–173). Berlin: Springer.
- Schaerf, A. (1999). A survey of automated timetabling. *Artificial Intelligence Review*, 13(2), 87–127.
- Schimmelpfeng, K., & Helber, S. (2007). Application of a real-world university-course timetabling model solved by integer programming. *OR Spectrum*, 29, 783–803.