

TIMO BERTHOLD^{*}
STEFAN HEINZ^{*}
MARCO E. LÜBBECKE²
ROLF H. MÖHRING³
JENS SCHULZ³

A Constraint Integer Programming Approach for Resource-Constrained Project Scheduling

^{*} Supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

² Technische Universität Darmstadt, Fachbereich Mathematik, Dolivostr. 15,64293 Darmstadt, Germany

³ Technische Universität Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany

A Constraint Integer Programming Approach for Resource-Constrained Project Scheduling

Timo Berthold^{1,*}, Stefan Heinz^{1,*}, Marco E. Lübbecke², Rolf H. Möhring³,
and Jens Schulz³

¹ Zuse Institute Berlin, Takustr. 7, 14195 Berlin, Germany
{berthold,heinz}@zib.de

² Technische Universität Darmstadt, Fachbereich Mathematik, Dolivostr. 15,
64293 Darmstadt, Germany
luebbecke@mathematik.tu-darmstadt.de

³ Technische Universität Berlin, Institut für Mathematik, Straße des 17. Juni 136,
10623 Berlin, Germany
{moehring,jschulz}@math.tu-berlin.de

Abstract. We propose a hybrid approach for solving the resource-constrained project scheduling problem which is an extremely hard to solve combinatorial optimization problem of practical relevance. Jobs have to be scheduled on (renewable) resources subject to precedence constraints such that the resource capacities are never exceeded and the latest completion time of all jobs is minimized.

The problem has challenged researchers from different communities, such as integer programming (IP), constraint programming (CP), and satisfiability testing (SAT). Still, there are instances with 60 jobs which have not been solved for many years. The currently best known approach, LAZYFD, is a hybrid between CP and SAT techniques.

In this paper we propose an even stronger hybridization by integrating all the three areas, IP, CP, and SAT, into a single branch-and-bound scheme. We show that lower bounds from the linear relaxation of the IP formulation and conflict analysis are key ingredients for pruning the search tree. First computational experiments show very promising results. For five instances of the well-known PSPLIB we report an improvement of lower bounds. Our implementation is generic, thus it can be potentially applied to similar problems as well.

1 Introduction

The resource-constrained project scheduling problem (RCPSP) is not only theoretically hard [4] but consistently resists computational attempts to obtain solutions of proven high quality even for instances of moderate size. As the problem is of high practical relevance, it is an ideal playground for different optimization communities, such as integer programming (IP), constraint programming (CP), and satisfiability testing (SAT), which lead to a variety of publications, see [5].

* Supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

The three areas each come with their own strengths to reduce the size of the search space. Integer programming solvers build on lower bounds obtained from linear relaxations. Relaxation can often be considerably strengthened by additional valid inequalities (*cuts*), which spawned the rich theory of polyhedral combinatorics. Constraint programming techniques cleverly learn about logical implications (between variable settings) which are used to strengthen the bounds on variables (*domain propagation*). Moreover, the constraints in a CP model are usually much more expressive than in the IP world. Satisfiability testing, or SAT for short, actually draws from *unsatisfiable* or conflicting structures which helps to quickly finding reasons for and excluding infeasible parts of the search space. The RCPSP offers footholds to attacks from all three fields but no single one alone has been able to crack the problem. So it is not surprising that the currently best known approach [10] is a hybrid between two areas, CP and SAT. Conceptually, it is the logical next step to integrate the IP world as well. It is the purpose of this study to evaluate the potential of such a hybrid and to give a proof-of-concept.

Our contribution. Following the constraint integer programming (CIP) paradigm as realized in SCIP [1, 11], we integrate the three techniques into a single branch-and-bound tree. We present a CP approach, enhanced by lower bounds from the linear programming (LP) relaxation, supported by the SCIP intern conflict analysis and a problem specific heuristic. We evaluate the usefulness of LP relaxation and conflict analysis in order to solve scheduling problems from the well known PSPLIB [9]. CP’s global *cumulative* constraint is an essential part of our model, and one contribution of our work is to make this constraint generically available within the CIP solver SCIP.

In our preliminary computational experiments it turns out that already a basic implementation is competitive with the state-of-the-art. It is remarkable that this holds for both, upper *and* lower bounds, the respective best known of which were not obtained with a single approach. In fact, besides meeting upper bounds which were found only very recently [10], independently of our work, we improve on several best known lower bounds of instances of the PSPLIB.

Related work. For an overview on models and techniques for solving the RCPSP we refer to the recent survey of [5]. Several works on scheduling problems already combine solving techniques in hybrid approaches. For the best current results on instances of PSPLIB, we refer to [10], where a constraint programming approach is supported by lazily creating a SAT model during the branch-and-bound process by which new constraints, so called *no-goods*, are generated.

2 Problem description

In the RCPSP we are given a set \mathcal{J} of non-preemptable jobs and a set \mathcal{R} of renewable resources. Each resource $k \in \mathcal{R}$ has bounded capacity $R_k \in \mathbb{N}$. Every job j has a processing time $p_j \in \mathbb{N}$ and resource demands $r_{jk} \in \mathbb{N}$ of each

resource $k \in \mathcal{R}$. The starting time S_j of a job is constrained by its predecessors that are given by a precedence graph $D = (V, A)$ with $V \subseteq \mathcal{J}$. An arc $(i, j) \in A$ represents a precedence relationship, i.e., job i must be finished before job j starts. The goal is to schedule all jobs with respect to resource and precedence constraints, such that the latest completion time of all jobs is minimized.

The RCPSP can be modeled easily as a constraint program using the global **cumulative** constraint [3] which enforces that at each point in time, the cumulated demand of the set of jobs running at that point, does not exceed the given capacities. Given a vector \mathbf{S} of start time variables S_j for each job j , the RCPSP can be modeled as follows:

$$\begin{aligned} \min \quad & \max_{j \in \mathcal{J}} S_j + p_j \\ \text{subject to} \quad & S_i + p_i \leq S_j && \forall (i, j) \in A \\ & \text{cumulative}(\mathbf{S}, \mathbf{p}, \mathbf{r}_{\cdot k}, R_k) && \forall k \in \mathcal{R} \end{aligned} \quad (1)$$

3 Linear programming relaxation and conflict analysis

For the implementation we use the CIP solver SCIP which performs a complete search in a branch-and-bound manner. The question to answer is how strongly conflict analysis and LP techniques are involved in the solving process by pruning the search tree. Therefore a first version of separation and conflict analysis methods are implemented for the cumulative constraint.

As IP model we use the formulation of [8] with binary start time variables. In the **cumulative** constraint we generate knapsack constraints [1] from the capacity cuts. Propagation of variable bounds and repropagations of bound changes are left to the solver SCIP. For the **cumulative** constraint bounds are updated according to the concept of *core-times* [6]. The core-time of a job is defined by the interval $[ub_j, lb_j + p_j]$. A jobs lower bound can be updated from lb_j to lb_j^* if its demand plus the demands of the cores exceed the resource capacity in certain time intervals. An explanation of this bound change is given by the set of jobs that have a core during this interval. More formally, let $\mathcal{C} \subset \mathcal{J}$ be the set of jobs whose core is non-empty, i.e., $ub_j < lb_j + p_j$ holds for $j \in \mathcal{C}$. The delivered explanation is the local lower bound of job j itself and the local lower and upper bounds of all jobs $k \in \{i \in \mathcal{C} : ub_i < lb_i^* \text{ and } lb_i + p_i > lb_j\}$.

This poses the interesting still open question whether it is NP-hard to find a minimum set of jobs from which the bound change can be derived.

To speed up the propagation process, we filter from the **cumulative** constraints, all pairs of jobs that cannot be executed in parallel and propagate them in a global **disjunctive** bounds constraint. This one propagates and checks the constraints in a more efficient manner and can separate further cuts based on forbidden sets. To get tight primal bounds, we apply a primal heuristic that is based on a fast list scheduling algorithm [7]. If an LP solution is available the list of jobs is sorted according to the start times of the jobs, otherwise by weighted local bounds, and α -points [7]. Furthermore, we apply a *justification* improve-

Table 1. Summary of the computational results. Detailed results are given in the appendix.

Setting	480 instances with 30 jobs							480 instances with 60 jobs						
				Nodes		Time in [s]					Nodes		Time in [s]	
	opt	best	wor.	total(k)	geom	total(k)	geom	opt	best	wor.	total(k)	geom	total(k)	geom
default	460	476	4	3 513	173.2	93.0	7.8	385	395	85	34 104	364.3	350.9	27.3
noconflict	436	467	13	8 665	246.6	175.0	11.6	381	390	90	38 099	381.8	362.9	28.3
norelax	454	467	13	7 444	194.0	106.8	6.5	384	390	90	127 684	591.2	355.8	26.1
none	446	465	15	9 356	217.5	135.5	7.7	382	389	91	126 714	599.3	364.8	26.9
bestset	460	476	4	–	–	–	–	391	401	79	–	–	–	–
LAZYFD	480	480	0	–	–	–	–	429	429	51	–	–	–	–

ment heuristic as described in [12] whenever a better solution was found. We use hybrid branching [2] only on integer variables.

4 Computational results

In this section, we analyze the impact of the two features LP relaxation and conflict analysis for the RCPSP using the test sets of the PSPLIB [9]. Due to the lack of space we restrict ourselves mainly to the test sets containing 30 and 60 jobs. For instances with 120 jobs we report improved lower bounds.

All computations were obtained on Intel Xeon Core 2.66 GHz computers (in 64 bit mode) with 4 MB cache, running Linux, and 8 GB of main memory. We used SCIP [11] version 1.2.0.6 and integrated CPLEX release version 12.10 as underlying LP solver. A time limit of one hour was enforced for each instance.

Table 1 presents the results for different settings which differ by disabled features. The setting “norelax” does not take advantage of the LP relaxation, “noconflict” avoids conflict analysis, “none” stands for disabling both these features whereas “default” enables both. The settings “bestset” is the best of the previous four settings for each instance and the last line reports the results for the solver LAZYFD. We compare for how many instances optimality (“opt”) was proven, the best known primal solution (“best”) was found, and the primal solution was worse (“wor.”) than the best known. Besides that we state total time and number of branch-and-bound nodes over all instances in the test set and the shifted geometric means⁴ (“geom”) over these two performance measures.

First of all the results show that our approach is competitive to the current best known method [10]. We observe further, that using both features leads to a tremendous reduction of the search space. This does not directly transfer to the running time. From that point of view the relaxation seems to be more expensive as the conflict analysis. On the other hand, the relaxation prunes a greater portion of the search space compared to the reduction achieved by the

⁴ The shifted geometric mean of values t_1, \dots, t_n is defined as $(\prod(t_i + s))^{1/n} - s$ with shift s . We use a shift $s = 10$ for time and $s = 100$ for nodes in order to decrease the strong influence of the very easy instances in the mean values.