
Semidefinite Programming Relaxations in Timetabling (Abstract)

Edmund K. Burke · Jakub Mareček · Andrew J. Parkes

Keywords timetabling · bounded colouring · vertex colouring · graph colouring · semidefinite programming

Semidefinite programming has recently gained much attention as a powerful method for deriving both strong lower bounds and approximation algorithms in combinatorial optimisation. There have not been, however, any applications to timetabling. We show one reason to believe that this could well change, ultimately.

Definitions In linear programming (LP), the task is to optimise a linear combination $c^T x$ subject to linear constraints $Ax = b$, together with the constraint that each in vector x of n variables is non-negative. The non-negativity of x , $x \in (R^+)^n$, can be seen as a restriction of the variables to lie in the convex cone of the positive orthant. Using interior point methods, linear programming can be solved to any fixed precision in polynomial time. These methods also work for other symmetric convex cones.

Semidefinite programming (SDP, Bellman & Fan, 1963; Alizadeh, 1995; Wolkowicz, Saigal, & Vandenberghe, 2000) is a generalisation of linear programming, replacing the vector variable with a square symmetric matrix variable and the polyhedral symmetric convex cone of the positive orthant with the non-polyhedral symmetric convex cone of positive semidefinite matrices. Note that an $n \times n$ matrix, M , is positive semidefinite if and only if $y^T M y \geq 0$ for all $y \in \mathbb{R}^n$. As all scalar multiples of positive semidefinite matrices and convex combinations of pairs of positive semidefinite matrices are positive semidefinite, positive semidefinite matrices do form a convex cone in R^{n^2} . We denote $A \succeq B$ whenever $A - B$ is positive semidefinite, and use $\langle A, B \rangle$ for the inner product of matrices, which is $\sum_{i,j} A_{i,j} B_{j,i}$. Formally, semidefinite programming is the minimisation of $\langle C, X \rangle$ such that $\langle A_i, X \rangle = b_i \quad \forall i = 1 \dots m$ and $X \succeq 0$, where X is a (primal) square symmetric matrix variable, C and A_i are compatible symmetric matrices, m is the number of constraints, and $b \in \mathbb{R}^m$.

Let us now consider a simple model of timetabling, underlying integer programming decompositions (Burke, Mareček, Parkes, & Rudová, 2010), for instance. The input consists of identifiers of events V , distinct enrolments U (“curricula”), rooms R , and periods P , plus mapping $F : U \rightarrow 2^V \setminus \emptyset$ from “curricula” to non-empty sets of events. Conflict graph $G = (V, E)$ is given

Automated Scheduling, Optimisation and Planning Group, School of Computer Science, The University of Nottingham, Nottingham NG8 1BB, UK. E-mail: {ekb,jxm,ajp}@cs.nott.ac.uk. For supplementary material, please see: <http://cs.nott.ac.uk/~jxm/>

by F , where events $F(u)$ is a clique in G for all $u \in U$. The “core” decision variables are

$$Z_{p,v} = \begin{cases} 1 & \text{event } v \text{ is taught at period } p \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

which are subject to linear constraints

$$\forall v \in V \quad \sum_{p \in P} Z_{p,v} = 1 \quad (2)$$

$$\forall p \in P \quad \forall u \in U \quad \sum_{v \in F_u} Z_{p,v} \leq 1 \quad (3)$$

$$\forall p \in P \quad \sum_{v \in V} Z_{p,v} \leq |R| \quad (4)$$

Notice that there is only a single mention (4) of rooms, which makes the colouring of the conflict graph $|R|$ -bounded. This means the cardinality of each colour class or the number of uses of each colour is at most $|R|$. Depending on the tightness of the $|R|$ -bound, the chromatic number alone is not necessarily a good lower bound.

Related Work There are a number of ways to bound the chromatic number of a graph using SDP. Informally, the point is that a parameter of the graph, denoted theta, is at least as large as the clique number and no more than the chromatic number, yet is computable in polynomial time using SDP. The known theta-like bounds for unbounded colouring form a hierarchy (Szegedy, 1994):

$$\alpha(G) \leq \vartheta_{1/2}(G) \leq \vartheta(G) \leq \vartheta_2(G) \leq \chi(\bar{G}), \text{ or } \omega(G) \leq \vartheta_{1/2}(\bar{G}) \leq \vartheta(\bar{G}) \leq \vartheta_2(\bar{G}) \leq \chi(G),$$

where α is the size of the largest independent set, ω is the size of the largest clique, χ is the chromatic number, $\vartheta_{1/2}$ is the vector chromatic number (Karger, Motwani, & Sudan, 1998), ϑ is the strict vector chromatic number (Karger et al., 1998), ϑ_2 is the strong vector chromatic number (Kleinberg & Goemans, 1998), and bar indicates complementation. For the corresponding vector programming and semidefinite programming formulations, please consult the literature (Szegedy, 1994). In theory, all could be extended to bounded graph colouring, but none has been so far, up to the best of our knowledge.

In terms of applications, the celebrated SDP relaxation of the maximum cut problem (MAX-CUT, Goemans & Williamson, 1995) has been adapted to scheduling workload on two machines (Skutella, 2001; Yang, Ye, & Zhang, 2003) and home-away patterns in sports scheduling (Suzuka, Miyashiro, Yoshise, & Matsui, 2007). The techniques of “vector lifting” and “matrix lifting” have been applied in signal decoding in multi-antenna systems (Mobasher & Khandani, 2007; Mobasher, Taherzadeh, Sotirov, & Khandani, 2007). All of the above can be thought of as rank-minimisation matrix completion problems (Fazel, Hindi, & Boyd, 2004), whose applications range from signal processing to statistics and system theory. We are now aware, however, of any applications to timetabling.

Bounding the Bounded Chromatic Number by SDP A clear application of semidefinite programming is in the detection of infeasibility in timetabling (2–4). The infeasibility test is given by lower bounding the $|R|$ -bounded chromatic number of the conflict graph and comparing it against $|P|$, the number of periods available. Here we follow the method and notation of (Dukanovic & Rendl, 2007), briefly reported also in PATAT 2004 (Dukanovic & Rendl, 2004). The underlying matrix variable M is:

$$M_{u,v} = \begin{cases} 1 & \text{if } u \text{ and } v \text{ are of the same colour} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

If we define $Y = tM$, we obtain legal colouring for integral t and $Y \in \{0, t\}$. In computing theta, these integrality constraints are dropped, resulting in an instance of SDP.

Table 1: An illustration of the effects of bounding the $|\mathbf{R}|$ -bounded chromatic number of the instance sta-f-83: Column $\chi^{|\mathbf{R}|}$ lists the $|\mathbf{R}|$ -bounded chromatic number obtained using integer linear programming, within time listed under “ $\chi^{|\mathbf{R}|}$ Runtime” in seconds. Column $\vartheta^{|\mathbf{R}|}$ lists the bounds obtained using semidefinite programming and rounding up, within time listed under “ $\vartheta^{|\mathbf{R}|}$ Runtime” in seconds. Column $|V|/|\mathbf{R}|$ lists the lower bound on the colours obtained by simple counting arguments and rounding up. Dash denotes the omission of the $|\mathbf{R}|$ -bounding constraint, giving the standard theta function instead of $\vartheta^{|\mathbf{R}|}$.

$ \mathbf{R} $	$\chi^{ \mathbf{R} }$	$\chi^{ \mathbf{R} }$ Runtime	$\vartheta^{ \mathbf{R} }$	$\vartheta^{ \mathbf{R} }$ Runtime	$ V / \mathbf{R} $
1	47	0.09	47	3.46	47
2	26	2.88	26	2.92	24
3	20	2.67	20	3.34	16
4	16	7.22	16	3.70	12
5	14	11.10	14	3.24	10
6	13	2.67	13	3.12	8
7	12	8.77	12	3.26	7
8	11	2.89	11	3.40	6
9	11	3.39	11	3.14	6
47	11	0.35	11	3.92	1
—	11	0.34	11	3.45	—

The theta relaxation can be modified to provide a bound on the bounded chromatic number by the addition of linear inequalities. In bounded colouring, we expect $\sum_u M_{uv} \leq |\mathbf{R}| \quad \forall v \in V$. This gives us the following SDP:

$$\vartheta^{|\mathbf{R}|}(\overline{G}) = \min t \quad (6)$$

$$\text{s. t. :} \quad \forall v \in V \quad Y_{vv} = t \quad (7)$$

$$\forall \{u, v\} \in E \quad Y_{uv} = 0 \quad (8)$$

$$\forall v \in V \quad \sum_u Y_{uv} \leq t|\mathbf{R}| \quad (9)$$

$$Y - J \succeq 0 \quad (10)$$

where J is the all-ones matrix. A closely related bound can be derived using the matrix lifting operator $M_+(K)$ of Lovász and Schrijver (1991).

Numerical Experiments As a concrete example, we consider a small conflict graph from a standard benchmark problem. Specifically, we take the instance ”sta-f-83” from the Toronto examination timetabling benchmarks ¹. There are 139 events, but the conflict graph has three connected components of 30, 47 and 62 vertices. Here, we use the 47-vertex component to study semidefinite programs produced by YALMIP (Löfberg, 2004) and solved using SeDuMi 1.21 (Sturm, 1999) and MathWorks Matlab R2009a on an Intel Core Duo P8600 at 2.4 GHz with 2 GB of RAM. For comparison, the bounded chromatic numbers are also provided. These were obtained using the most straightforward integer linear programming formulation solved using the defaults of ILOG CPLEX 12.10 on the same machine. Results are given in Table 1. Firstly, note that $|\mathbf{R}| = 1$ gives precisely the number of nodes, as would be expected. Secondly, note that $\vartheta^{|\mathbf{R}|}$ is generally much tighter than the lower bound $|V|/|\mathbf{R}|$ obtained by simple counting arguments. Accidentally, $\vartheta^{|\mathbf{R}|}$ lower bounds actually happen to match the optima in this particular instance. For example, at $|\mathbf{R}| = 5$, counting cannot rule out a 10-colouring, but the SDP bound shows that at least 14 colours are required. A 14-colouring together with a certificate of its optimality can be obtained using CPLEX, but not in polynomial time. As far as we know, SDP relaxations are the only way to get such information in polynomial time.

¹ See <ftp://ftp.mie.utoronto.ca/pub/carter/testprob/> and <http://www.cs.nott.ac.uk/~rxq/data.htm>

We have also tried $|R|$ -bounded modifications of the extensions to theta as given in (Dukanovic & Rendl, 2007): $\vartheta^{R|+}$ by keeping the $Y \geq 0$ constraints, and $\vartheta^{R|+\Delta}$ keeping the $Y \geq 0$ constraints and also adding triangle inequality constraints. These, however, slow down the solver and do not improve the bound on the tested instances. Also, theta can also be formulated on the complement graph, and this might be useful when the edge density is high, but we have not yet explored $|R|$ -bounded versions.

Future Work In the SDP relaxation of bounded graph colouring above, the colour assignment was not represented directly, but only in terms of the “same-time” classes of equivalence of nodes assigned the same colour. This makes it naturally invariant under permutation of the colours. This is sufficient for bounded colouring, but many objectives in timetabling refer to time-based patterns of activities, e.g. whether events should be on the same day or not. These are not invariant under “colour permutations” and so the “same-time” representation is no longer sufficient. For example, in lower bounding the **Surface** component of integer programming decompositions (Burke et al., 2010), i.e. the assignment of events to periods, including all the respective terms of the objective function, we presumably need to re-introduce some matrix variable mapping events to timeslots as in **Surface** (1). The matrix variable will need to be constrained so that there is only a single event in each roomslot. This gives a constraint on the rank of the matrix variable, and this can then be expressed in SDP. This can also be thought of as an application of matrix-lifting operator $M_+(K)$ of Lovász and Schrijver (Lovász & Schrijver, 1991). Work in this direction is in progress.

Conclusions The aim of this abstract was not to present a practical method for bounding the optima in timetabling problems, yet. Indeed, SDP solvers are less well-developed than LP solvers, in general. Current interior point methods for semidefinite programming are rather slow, albeit running in time polynomial in the dimensions of the instance for any fixed precision. Our hope, however, is that SDP solvers will improve significantly in the future. There is some evidence that this could happen (Monteiro, 2003). The nascent bundle (Helmberg & Rendl, 2000; Helmberg, 2003) and augmented Lagrangian methods (Burer & Vandenbussche, 2006) are particularly promising, as they seem to be able to cope with thousands of vertices in the conflict graph.

Notwithstanding the caveat above, SDP provides some of the strongest known relaxations in timetabling. An extension of theta to bounded graph colouring gives a useful lower bound on the number of periods required in the timetable, considering the conflict graph and the number of rooms. More complex relaxations seem to allow for the optimisation over the assignments of events to periods and rooms as well.

References

- Alizadeh, F. (1995). Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM J. Optim.*, 5(1), 13–51.
- Bellman, R., & Fan, K. (1963). On systems of linear inequalities in Hermitian matrix variables. In *Proc. Sympos. Pure Math., Vol. VII* (pp. 1–11). Providence, R.I.: Amer. Math. Soc.
- Burer, S., & Vandenbussche, D. (2006). Solving lift-and-project relaxations of binary integer programs. *SIAM J. Optim.*, 16(3), 726–750 (electronic).
- Burke, E. K., Mareček, J., Parkes, A. J., & Rudová, H. (2010). Decomposition, reformulation, and diving in university course timetabling. *Comput. Oper. Res.*, 37(1), 582–597.
- Dukanovic, I., & Rendl, F. (2004). Combinatorial tricks and Lovasz theta function applied to graph coloring. In *Proc. PATAT 2004* (p. 479).
- Dukanovic, I., & Rendl, F. (2007). Semidefinite programming relaxations for graph coloring and maximal clique problems. *Math. Program.*, 109(2-3, Ser. B), 345–365.
- Fazel, M., Hindi, H., & Boyd, S. (2004). Rank minimization and applications in system theory. In *American control conference* (pp. 3273–3278). AACC.

- Goemans, M. X., & Williamson, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6), 1115–1145.
- Helmberg, C. (2003). Numerical evaluation of SBmethod. *Math. Program.*, 95(2, Ser. B), 381–406. (Computational semidefinite and second order cone programming: the state of the art)
- Helmberg, C., & Rendl, F. (2000). A spectral bundle method for semidefinite programming. *SIAM J. Optim.*, 10(3), 673–696.
- Karger, D., Motwani, R., & Sudan, M. (1998). Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2), 246–265.
- Kleinberg, J., & Goemans, M. X. (1998). The Lovász theta function and a semidefinite programming relaxation of vertex cover. *SIAM J. Discrete Math.*, 11(2), 196–204.
- Löfberg, J. (2004). Yalmip : A toolbox for modeling and optimization in MATLAB. In *Proc. of CACSD*. Taipei, Taiwan.
- Lovász, L., & Schrijver, A. (1991). Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optim.*, 1(2), 166–190.
- Mobasher, A., & Khandani, A. K. (2007). Matrix-lifting semi-definite programming for decoding in multiple antenna systems. *CoRR*, abs/0709.1674.
- Mobasher, A., Taherzadeh, M., Sotirov, R., & Khandani, A. K. (2007). A near-maximum-likelihood decoding algorithm for MIMO systems based on semi-definite programming. *IEEE Trans. Inform. Theory*, 53(11), 3869–3886.
- Monteiro, R. D. C. (2003). First- and second-order methods for semidefinite programming. *Math. Program.*, 97(1-2, Ser. B), 209–244.
- Skutella, M. (2001). Convex quadratic and semidefinite programming relaxations in scheduling. *J. ACM*, 48(2), 206–242.
- Sturm, J. F. (1999). Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.*, 11/12(1-4), 625–653.
- Suzuka, A., Miyashiro, R., Yoshise, A., & Matsui, T. (2007). The home-away assignment problems and break minimization/maximization problems in sports scheduling. *Pac. J. Optim.*, 3(1), 113–133.
- Szegedy, M. (1994). A note on the theta number of Lovász and the generalized Delsarte bound. In *Sfcs '94: Proceedings of the 35th annual symposium on foundations of computer science* (pp. 36–39). Washington, DC, USA: IEEE Computer Society.
- Wolkowicz, H., Saigal, R., & Vandenberghe, L. (Eds.). (2000). *Handbook of semidefinite programming*. Boston, MA: Kluwer Academic Publishers. (Theory, algorithms, and applications)
- Yang, H., Ye, Y., & Zhang, J. (2003). An approximation algorithm for scheduling two parallel machines with capacity constraints. *Discrete Appl. Math.*, 130(3), 449–467.