

GRAPH COLOURING PROBLEMS AND THEIR APPLICATIONS IN SCHEDULING

Dániel MARX

Department of Computer Science and Information Theory
Budapest University of Technology and Economics
H-1521 Budapest, Hungary
e-mail: dmarx@cs.bme.hu

Received: Dec. 12, 2003

Abstract

Graph colouring and its generalizations are useful tools in modelling a wide variety of scheduling and assignment problems. In this paper we review several variants of graph colouring, such as precolouring extension, list colouring, multicolouring, minimum sum colouring, and discuss their applications in scheduling.

Keywords: scheduling, colouring, precolouring extension, list colouring, multicolouring, minimum sum colouring.

1. Introduction

A k -colouring of graph G is an assignment of integers $\{1, 2, \dots, k\}$ (the colours) to the vertices of G in such a way that neighbours receive different integers. The *chromatic number* of G is the smallest k such that G has a k -colouring.

There are several interesting practical problems that can be modelled by graph colouring. Our basic example is the following. Assume that we have to schedule a set of interfering jobs, it has to be determined when each job is executed. Let G be the *conflict graph* of the jobs: the vertices of the graph corresponds to the jobs, two vertices are connected by an edge if the corresponding two jobs cannot be executed at the same time (for example, they use a shared resource or interfere in some other ways). The colours correspond to the available time slots, every job requires one time slot. There is a one-to-one correspondence between the feasible schedulings of the jobs and the colourings of the graph: vertex v receives colour i if and only if the corresponding job is executed in time slot i . Clearly, the graph has a k -colouring if and only if the jobs can be done in k time slots such that interfering jobs are not executed simultaneously. Therefore the chromatic number of the graph equals the minimum *makespan* of the scheduling problem, the minimum time required to finish the jobs.

Unfortunately, determining the chromatic number of a graph is an **NP**-hard problem, hence we cannot expect to solve it efficiently for large graphs. Therefore when we model a scheduling or assignment problem with graph colouring, then there are two things to consider. First, it might happen that the graphs arising in our

application have a special structure that makes colouring easier. Moreover, if there is no hope for an efficient (polynomial time) algorithm to the colouring problem that always finds an optimum solution, then we have to content ourselves with an *approximation algorithm* that is not optimal, but has some performance guarantee on the quality of the produced solution. For a minimization problem, we say that an algorithm is a c -approximation algorithm if it always produces a solution whose cost is at most c times the optimum.

We conclude this section by citing some specific applications of graph colouring where the conflict graph has a special structure, and efficient optimal or approximation algorithm can be given for the problem.

Aircraft scheduling. Assume that we have k aircrafts, and we have to assign them to n flights, where the i th flight is during the time interval (a_i, b_i) . Clearly, if two flights overlap, then we cannot assign the same aircraft to both flights. The vertices of the conflict graph correspond to the flights, two vertices are connected if the corresponding time intervals overlap. Therefore the conflict graph is an *interval graph*, which can be coloured optimally in polynomial time.

Biprocessor tasks. Assume that we have a set of processors (machines) and a set of tasks, each task has to be executed on two preassigned processors simultaneously. A processor cannot work on two jobs at the same time. For example, such biprocessor tasks arise when we want to schedule file transfers between processors [1] or in the case of mutual diagnostic testing of processors [2].

Consider the graph whose vertices correspond to the processors, and if there is a task that has to be executed on processors i and j , then we add an edge between the two corresponding vertices. Now the scheduling problem can be modelled as an *edge colouring* of this graph: we have to assign colours to the edges in such a way that every colour appears at most once at a vertex.

Edge colouring is **NP**-hard [3], but there are good approximation algorithms. The maximum degree Δ of the graph is an obvious lower bound on the number of colours needed to colour the edges of the graph. On the other hand, if there are no multiple edges in the graph (there are no two tasks that require the same two processors), then Vizing's Theorem gives an efficient method for obtaining a $(\Delta + 1)$ -edge colouring. If multiple edges are allowed, then the algorithm of [4] gives a 1.1-approximate solution.

Frequency assignment. Assume that we have a number of radio stations, identified by x and y coordinates in the plane. We have to assign a frequency to each station, but due to interferences, stations that are 'close' to each other have to receive different frequencies. Such problems arise in frequency assignment of base stations in cellular phone networks.

At first sight, one might think that the conflict graph is planar in this problem, and the Four Colour Theorem can be used, but it is not true: if there are lots of stations in small region, then they are all close to each other, therefore they form a large

clique in the conflict graph. Instead, the conflict graph is a *unit disk graph*, where each vertex corresponds to a disk in the plane with unit diameter, and two vertices are connected if and only if the corresponding disks intersect. A 3-approximation algorithm for colouring unit disk graphs is given in [5], yielding a 3-approximation for the frequency assignment problem.

2. Multicolouring

A natural generalization of the basic setup introduced in Section 1 is to consider jobs that require more than one time slots. In the *multicolouring* problem each vertex v has a *demand* $x(v)$, and we have to assign a set of $x(v)$ colours to each vertex v such that neighbours receive disjoint sets of colours. Multicolouring can be used to model the scheduling of jobs with different time requirements: the set of colours assigned to vertex v corresponds to the $x(v)$ time slots when we work on the job.

There are two main variants of multicolouring. In *non-preemptive* multicolouring the set of colours assigned to a vertex has to be a continuous interval of colours. This reflects the requirement that the jobs cannot be interrupted, they have to receive a continuous time window. On the other hand, in *preemptive* multicolouring we assume that the jobs can be interrupted, hence the set of colours assigned to a vertex can be arbitrary, it does not have to be continuous.

In [6] the frequency assignment problem discussed in Section 1 is generalized, we have to assign a predefined number of frequencies to each base station. The authors adopt the hexagon model, which means that the conflict graph is a subgraph of the triangular lattice. A $\frac{4}{3}$ -approximation algorithm is presented in [6] for minimizing the number of different frequencies assigned.

3. Precolouring Extension

In certain scheduling problems we do not have full control over the schedule, the assignments of certain jobs are already decided. In this case some of the vertices of the conflict graph have a preassigned colour, and we have to solve the *precolouring extension* problem: extend the colouring of these vertices to the whole graph, using the minimum number of colours.

BIRÓ, HUJTER and TUZA [7, 8, 9] started a systematic study of precolouring extension. In [7], the aircraft scheduling problem discussed in Section 1 is extended. There is a maintenance period for each aircraft, during which it cannot fly. We can model these maintenance periods by adding a ‘dummy’ flight for the maintenance period of each aircraft, and requiring that the maintenance period of the i th aircraft is assigned to the i th aircraft. Therefore we have to solve the precolouring extension problem on the conflict graph, which is an interval graph. It is shown in [7] that the precolouring extension problem is **NP**-complete for interval graphs, but it can be

solved in polynomial time if every colour is used only once in the precolouring, that is, if every aircraft has only one maintenance interval (the latter result is generalized to chordal graphs in [10]).

4. List Colouring

In the list colouring problem each vertex v has a list of available colours, and we have to find a colouring where the colour of each vertex is taken from its list of available colours. List colouring can be used to model situations where a job can be processed only in certain time slots, or if it can be processed only by certain machines.

Using standard dynamic programming techniques, list colouring can be solved in polynomial time on trees and partial k -trees [11]. By combining dynamic programming with a clever use matching, list colouring can be solved on the edges of trees as well [12].

The multicolouring concept introduced in Section 2 can be applied for list colourings as well: each vertex has an integer demand $x(v)$, and vertex v has to receive a set of $x(v)$ colours from its list of colours. The algorithm for list colouring trees and partial k -trees does not generalize for the multicolouring case, as the problem is **NP**-complete already for binary trees [13]. On the other hand, list *edge* multicolouring can be solved in polynomial time on trees: using standard techniques, the good characterization theorem of MARCOTTE and SEYMOUR [14] can be turned into a polynomial time algorithm. This result is generalized in [15] to a slightly more general class of graphs, that includes odd cycles. Moreover, a randomized algorithm is given for an even more general class of graphs, including even cycles.

5. Minimum Sum Colouring

Besides minimizing the makespan, another well-studied goal in scheduling theory is to minimize the sum of completion times of the jobs, which is the same as minimizing the average completion time. The corresponding colouring problem is *minimum sum colouring*, introduced in [16]: we are looking for a colouring of the conflict graph such that the sum of the colours assigned to the vertices is minimal.

Apart from trees, partial k -trees, and edges of trees, minimum sum colouring is **NP**-hard on most classes of graphs. On the other hand, it turns out that the sum of the colouring is easier to approximate than the makespan (see e.g. [17, 18] for approximation results). The reason for this is that the sum of the colouring and the makespan of the colouring behave very differently when a small part of the graph is recoloured. If we recolour a small part of the graph, then this change has only a small effect on the sum of the colouring, but it can change the makespan significantly.

The multicolouring version of the problem can be used to model arbitrary length jobs. Since we want to minimize the sum of the completion times, the objective function of the colouring problem has to be defined as follows. The *finish time* of a vertex is the largest colour assigned to it, and the sum of a colouring is the sum of the finish times of the vertices. It is clear that the sum of the finish times in a multicolouring is equal to the sum of completion times in the corresponding schedule. This variant of multicolouring was introduced in [19], where approximation algorithms are given for various classes of graphs. The preemptive and non-preemptive versions of the problem can have very different complexity: while the non-preemptive version can be solved in polynomial time for trees [20], the preemptive version is **NP**-hard for binary trees [13], but has a polynomial time approximation scheme [20]. In [21] polynomial time approximation schemes are given for partial k -trees and planar graphs as well. Unlike minimum sum colouring, the multicolouring version of the problem is **NP**-hard on the edges of trees. However, in this case the problem admits a polynomial time approximation scheme [22].

6. Acknowledgements

Research was supported by OTKA grants 44733, 42559 and 42706.

References

- [1] COFFMAN, E. G. – GAREY, JR. M. R. – JOHNSON, D. S. – LAPAUGH, A. S., Scheduling File Transfers, *SIAM J. Comput.*, **14** (3) (1985), pp. 744–780.
- [2] HOOGEVEEN, J. A. – VAN DE VELDE, S. L. – VELTMAN, B., Complexity of Scheduling Multiprocessor Tasks with Prespecified Processor Allocations, *Discrete Appl. Math.*, **55** (3) (1994), 259–272.
- [3] HOLYER, I., The NP-Completeness of Edge-Coloring, *SIAM J. Comput.*, **10** (4) (1981), pp. 718–720.
- [4] NISHIZEKI, T. – KASHIWAGI, K., On the 1.1 Edge-Coloring of Multigraphs, *SIAM J. Discrete Math.*, **3** (3) (1990), pp. 391–410.
- [5] GRÄF, A. – STUMPF, M. – WEIßENFELS, G., On Coloring Unit Disk Graphs, *Algorithmica*, **20** (3) (1998), pp. 277–293.
- [6] NARAYANAN, L. – SHENDE, S. M., Static Frequency Assignment in Cellular Networks, *Algorithmica*, **29** (3) (2001) pp. 396–409.
- [7] BIRÓ, M. – HUJTER, M. – TUZA, ZS., Precoloring Extension. I. Interval Graphs, *Discrete Math.*, **100** (1–3) (1992), pp. 267–279.
- [8] HUJTER, M. – TUZA, ZS., Precoloring Extension. II. Graph Classes Related to Bipartite Graphs, *Acta Mathematica Universitatis Comenianae*, **62** (1) (1993), pp. 1–11.
- [9] HUJTER, M. – TUZA, ZS., Precoloring Extension. III. Classes of Perfect Graphs, *Combin. Probab. Comput.*, **5** (1) (1996), pp. 35–56.
- [10] MARX, D., *Precoloring Extension on Chordal Graphs*, 2003, submitted.
- [11] JANSEN, K. – SCHEFFLER, P., Generalized Coloring for Tree-Like Graphs, *Discrete Appl. Math.*, **75** (2) (1997), pp. 135–155.
- [12] GIARO, K. – KUBALE, M., Edge-Chromatic Sum of Trees and Bounded Cyclicity Graphs, *Inform. Process. Lett.*, **75** (1-2) (200), pp. 65–69.

- [13] MARX, D., The Complexity of Tree Multicolorings, in *Mathematical Foundations of Computer Science 2002 (Warsaw-Otwock)*, pp. 532–542. Springer, Berlin, 2002.
- [14] MARCOTTE, O. – SEYMOUR, P. D., Extending an Edge-Coloring, *J. Graph Theory*, **14** (5) (1990), pp. 565–573.
- [15] MARX, D., List Edge Multicoloring in Graphs with Few Cycles, *Inform. Process. Lett.*, **89** (2) (2004), pp. 85–90.
- [16] KUBICKA, E. – SCHWENK, A. J., An Introduction to Chromatic Sums, in *Proceedings of the ACM Computer Science Conf.*, pp. 15–21. Springer, Berlin, 1989.
- [17] BAR-NOY, A. – BELLARE, M. – HALLDÓRSSON, M. M. – SHACHNAI, H. – TAMIR, T., On Chromatic Sums and Distributed Resource Allocation, *Inform. and Comput.*, **140** (2) (1998), pp. 183–202.
- [18] GIARO, K. – JANCZEWSKI, R. – KUBALE, M. – MAI AFIEJSKI, M., A $27/26$ -Approximation Algorithm for the Chromatic Sum Coloring of Bipartite Graphs, in *Proceedings of APPROX 2002*, (2002), pp. 135–145.
- [19] BAR-NOY, A. – HALLDÓRSSON, M. M. – KORTSARZ, G. – SALMAN, R. – SHACHNAI, H., Sum Multicoloring of Graphs, *J. Algorithms*, **37** (2) (2000), pp. 422–450.
- [20] HALLDÓRSSON, M. M. – KORTSARZ, G. – PROSKUROWSKI, A. – SALMAN, R. – SHACHNAI, H. – TELLE, J. A., Multicoloring Trees, *Inform. and Comput.*, 180(2):113–129, 2003.
- [21] HALLDÓRSSON, M. M. – KORTSARZ, G., Tools for Multicoloring with Applications to Planar Graphs and Partial k -Trees, *J. Algorithms*, **42** (2) (2002), pp. 334–366.
- [22] MARX, D., Minimum Sum Multicoloring on the Edges of Trees, *1st Workshop on Approximation and Online Algorithms (WAOA)*, Budapest, (2003) pp. 214–226, Lecture Notes in Comput. Sci., 2909, Springer, Berlin, 2004.