Expert Systems with Applications 42 (2015) 5491–5507

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

A grouping hyper-heuristic framework: Application on graph colouring

Anas Elhag *, Ender Özcan

ASAP Research Group, School of Computer Science, University of Nottingham, Jubilee Campus, NG8 1BB Nottingham, UK

ARTICLE INFO

Article history:
Available online 31 January 2015

Keywords:
Hyper-heuristics
Grouping problems
Graph colouring
Timetabling

ABSTRACT

Grouping problems are hard to solve combinatorial optimisation problems which require partitioning of objects into a minimum number of subsets while a given objective is simultaneously optimised. Selection hyper-heuristics are high level general purpose search methodologies that operate on a space formed by a set of low level heuristics rather than solutions. Most of the recently proposed selection hyper-heuristics are iterative and make use of two key methods which are employed successively; heuristic selection and move acceptance.

In this study, we present a novel generic selection hyper-heuristic framework containing a fixed set of reusable grouping low level heuristics and an unconventional move acceptance mechanism for solving grouping problems. This framework deals with one solution at a time at any given decision point during the search process. Also, a set of high quality solutions, capturing the trade-off between the number of groups and the additional objective for the given grouping problem, is maintained. The move acceptance mechanism embeds a local search approach which is capable of progressing improvements on those trade-off solutions.

The performance of different selection hyper-heuristics with various components under the proposed framework is investigated on graph colouring as a representative grouping problem. Then, the top performing hyper-heuristics are applied to a benchmark of examination timetabling instances. The empirical results indicate the effectiveness and generality of the proposed framework enabling grouping hyper-heuristics to achieve high quality solutions in both domains.

1. Introduction

The task of partitioning a large set of items into a collection of mutually disjoint subsets is a common task in a variety of real-world problems. In a *grouping problem*, the goal is to optimise a given objective (cost, penalty) while achieving the minimum number of subsets (groups). Hence, grouping problems can be formulated as a multi-criteria discrete combinatorial optimisation problem, considering that there is a trade-off between minimising cost and number of groups, as in graph colouring (Saha, Kumar, & Baboo, 2013), timetabling (Qu, Burke, McCollum, Merlot, & Lee, 2009) and packing (Falkenauer, 1998).

Two crucial components in the design of *grouping algorithms* for solving grouping problems are the candidate solution representation and neighbourhood/move operator(s). A redundant representation scheme which allows equivalent solutions yielding the same grouping creates a huge search space that might impair even the most powerful search algorithm. Many grouping approaches based on genetic algorithms (GAs) have been explored in the scientific literature providing various degrees of success (Falkenauer, 1998; Korkmaz, 2010). In previous studies, it has been observed that traditional operators are rather disruptive and, in many cases, counter productive, hence special operators that are tailored for grouping problems are needed.

There is a growing number of studies on more general and reusable search methodologies applicable to multiple problem domains than the existing specifically tailored solutions to a single problem. *Hyper-heuristics* are such high level search methodologies that search the space formed by low level heuristics, instead of solutions directly for solving hard problems (Burke et al., 2013). There are different types of hyper-heuristics. The focus of this study is *selection* type of high level search methods that mix and control a pre-defined set of low level perturbative heuristics (move operators) processing a single complete solution at each step under a single-point based search framework.

In this study, we describe a novel selection hyper-heuristic framework for grouping problems. The framework provides a set of tailored reusable low level grouping heuristics. In contrast to traditional selection hyper-heuristics that use a different set of low level heuristics provided for each different problem domain,

* Corresponding author.
E-mail addresses: axe@cs.nott.ac.uk (A. Elhag), exo@cs.nott.ac.uk (E. Özcan).

http://dx.doi.org/10.1016/j.eswa.2015.01.038
0957-4174/

in our proposed framework the set of low level heuristics is fixed and the same framework can be used for solving various grouping problems (see Section 2.3). This adds another level of generality when compared to generic selection hyper-heuristics.

Although a single-point based search framework is used, a set of solutions, capturing the trade-off between number of groups and some cost specific to the grouping problem in hand, is maintained. After a move operator creates a new solution from a given solution, move acceptance method is used to decide whether to accept or reject that resultant solution. Our framework contains an unconventional move acceptance component. This novel mechanism, designed for grouping problems, attempts to progress improvements in certain situations via a local search algorithm regardless of the acceptance/rejection decision.

Any component of the proposed framework can be implemented based on any appropriate grouping representation. However, in this study, we introduce and employ a modified (restricted) version of grouping representation, referred to as the *Group Encoding* (Falkenauer, 1998). We have investigated the performance of the framework using different selection hyper-heuristic components on a set of well known graph colouring benchmark instances.[1] Additionally, we applied the top hyper-heuristics without any modification to a benchmark of examination timetabling instances in order to examine the generality of the framework. The empirical results show that a learning selection hyper-heuristic developed using the framework turns out to be indeed sufficiently general and reusable. This hyper-heuristic either beats most of the previously proposed approaches tailored for the specific problem in hand or shows that it is highly competitive.

The paper is organised as follows. Section 2 provides an overview of grouping problems, different representation schemes for grouping problems and hyper-heuristics. The details of the proposed selection hyper-heuristic framework including all low level heuristics and different components are given in Section 3. The experimental design and results are discussed in Section 4, while the last section presents concluding remarks and future work.

2. Background

2.1. Grouping problems

Grouping problems are combinatorial optimisation problems in which a large group of n items, $U = \{x_1, x_2, x_3, \dots, x_n\}$, is to be divided into a collection of k $(2 \leqslant k \leqslant n-1)$ subgroups, u_i $(1 \leqslant i \leqslant k)$; such that each item $x \in U$ belongs to exactly one subgroup minimising a given objective (cost/penalty/fitness) and k. Different grouping problems have different constraints, and introduce different objective (cost) functions, as in graph colouring, timetabling, data clustering and packing (Falkenauer, 1998). In our formulation, we denote a cost function as a decomposable function, $f()$. For a subgroup u_i, the partial cost is denoted as $f(u_i)$, and for a complete solution $U_g = \{u_1, \dots, u_k\}$, $f(U_g)$ is the total cost.

$$\textit{minimise} \quad Z = \left(f(U_g) = \sum_{i=1}^{k} f(u_i), k \right) \tag{1}$$

$$\textit{subject to} \quad \bigcup_{i=1}^{k} u_i = U \tag{2}$$

$$u_i \bigcap u_j = \emptyset \quad \forall_{i,j} \text{ where } i \neq j \tag{3}$$

$$u_i \neq \emptyset \quad \forall_i \tag{4}$$

[1] <ftp://dimacs.rutgers.edu/pub/challenge/graph/benchmarks/color/>.

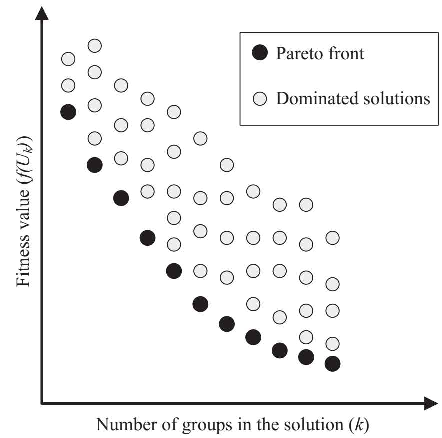

Fig. 1. The dominance concept in multi-objective optimisation.

In this study, we represent the grouping problem as a discrete two objective multi-criteria problem in which the goal is to optimise the two conflicting objectives in Eq. (1) above, namely the number of groups which can only take discrete values; and the cost which can take discrete or continuous values depending on the problem. Ideally, these two objectives should be *simultaneously* optimised, although they are clearly conflicting; i.e. a decrease in the number of groups k leads to an increase in the cost. In some cases, there might not be a single optimal solution. Instead, there could be multiple solutions with a trade-off from which a decision maker can choose. Those solutions are identified using the concept of *dominance* (Zitzler & Thiele, 1998) as illustrated in Fig. 1.

A solution x is considered to dominate another solution y, $(x \succ y)$ if, and only if, x is better than y in at least one objective, and x is not worse than y in any of the objectives. The set of the non-dominated solutions is known as the *Pareto optimal set*, and its image in the objective domain is known as the *Pareto optimal front*. This problem is different than a generic multi-objective problem where mostly, there is a region where the Pareto front is driven automatically via a multi-objective algorithm, however, in grouping problems the range of groups is fixed, hence the search methodology can focus on a single objective without ignoring the second one. We use some basic ideas from multi-objective optimisation, but the proposed approach is not a generic multi-objective algorithm as described in Section 3. For more on multi-objective optimisation, readers can refer to Zitzler and Thiele (1998), Coello, Lamont, and van Veldhuizen (2007), which is not the focus of this study.

2.1.1. Grouping representations

Almost all previously proposed grouping approaches are genetic algorithms utilising various encoding schemes for grouping problems. Those schemes can be classified as *fixed* and *variable* length representations. A fixed-length representation is based on an array of values associated in some way with each item in a set of objects, such as, *Group Numeric Encoding (GNE)* and *Permutation with Separators Encoding (PWS)* (Jones & Beltramo, 1991) which are widely used in the literature. Each location in the array is associated with an item and in GNE, the value at a location indicates the group that the item belongs to, whereas in PWS, that value represents the relative positions of the objects with respect to each other. However, many studies have concluded that such representations have some deficiencies. One of the crucial flaws is the redundancy in the representation. Different candidate solutions under a redundant representation could yield the same grouping of items. For example, assuming that we have 3 objects for