Operations Research Letters 45 (2017) 206–209

Contents lists available at ScienceDirect

Operations Research Letters

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

Limited memory Rank-1 Cuts for Vehicle Routing Problems

Diego Pecin [a,*], Artur Pessoa [b], Marcus Poggi [c], Eduardo Uchoa [b], Haroldo Santos [d]

[a] École Polytechnique de Montréal and GERAD, Canada
[b] Universidade Federal Fluminense, Departamento de Engenharia de Produção, Brazil
[c] Pontifícia Universidade Católica do Rio de Janeiro, Departamento de Informática, Brazil
[d] Universidade Federal de Ouro Preto, Departamento de Computação, Brazil

ARTICLE INFO

Article history:
Received 10 June 2016
Received in revised form
27 December 2016
Accepted 17 February 2017
Available online 24 February 2017

Keywords:
Set Partitioning
Polyhedral combinatorics
Branch-cut-and-price algorithms

ABSTRACT

Pecin et al. (2016) introduced a "limited memory" technique that allows an efficient use of Rank-1 cuts in the Set Partitioning Formulation of Vehicle Routing Problems, motivating a deeper investigation of those cuts. This work presents a computational polyhedral study that determines the best possible sets of multipliers for cuts with up to 5 rows. Experiments with CVRP instances show that the new multipliers lead to significantly improved dual bounds and contributes decisively for solving an open instance with 420 customers.

1. Introduction

The Vehicle Routing Problem (VRP) is among the most widely studied combinatorial optimization problems, due to its direct application in modern systems that distribute goods and services. Reflecting the large variety of conditions present in those systems, the VRP literature is spread into dozens of variants. For example, there are variants that consider vehicle capacities, time windows, multiple depots, heterogeneous fleet, pickups and deliveries, etc. The Set Partitioning Formulation (SPF) [3] can model most of those variants, the only assumption is that each customer should be visited once. Let $V = \{1, \ldots, n\}$ be the set of customers and let Ω be the set of all possible routes that respect the conditions of the considered variant. For each route $r \in \Omega$, let c_r be the cost of r, a_i^r be a binary coefficient that indicates whether r visits customer $i \in V$, and λ_r be a variable deciding if the route r is used or not. The formulation follows:

$$\text{Min:} \sum_{r\in\Omega} c_r \lambda_r \tag{1}$$

$$\text{S.t.:} \sum_{r\in\Omega} a_i^r \lambda_r = 1, \qquad \forall i \in V, \tag{2}$$

$$\lambda_r \in \{0, 1\}, \qquad \forall r \in \Omega. \tag{3}$$

Due to the huge number of variables, the linear relaxation of the SPF must be solved by column generation. The pricing subproblem consists in finding routes with minimum reduced cost, a problem that is often modeled as a shortest path with resource constraints and solved by *labeling algorithms* (see [22]).

The linear relaxation of the SPF is usually not strong enough to be the basis of state-of-the-art exact algorithms, at least for the most classical and competitive variants: the capacited VRP (CVRP) and the VRP with time windows (VRPTW). For that purpose, it must be strengthened with additional cuts. An algorithm that combines column and cut generation in a tree enumeration search is called a *branch-cut-and-price* (BCP) algorithm. According to the classification proposed in [21], *robust cuts* are those that do not change the structure of the pricing subproblem. In contrast, *non-robust cuts* change the pricing structure: each additional cut makes it harder, and so, too many cuts make it intractable. Robust cuts may be effective. In fact, some successful BCP algorithms [16,11,15,23] only use them. However, it seems that the potential for robust cuts in the classical variants is exhausted. In fact, no effective new family of robust cuts was found in the last decade. As a consequence, an important line in current research is finding effective non-robust cuts that are not very harmful to the labeling algorithms used in the pricing.

The Set Partitioning constraints (2) are the natural source of non-robust cuts. However, well-known families of cuts like clique or odd holes [1] make the pricing too expensive [24]. An important advance was the introduction of the Subset Row Cuts (SRCs) by Jepsen et al. [13]. The proposed SRCs are subfamilies of clique and

* Corresponding author.
E-mail addresses: diego.pecin@gerad.ca (D. Pecin), artur@producao.uff.br (A. Pessoa), poggi@inf.puc-rio.br (M. Poggi), uchoa@producao.uff.br (E. Uchoa), haroldo@iceb.ufop.br (H. Santos).

http://dx.doi.org/10.1016/j.orl.2017.02.006
0167-6377/

lifted odd holes defined over a small subset of the rows in (2). Those particular non-robust cuts can be better treated by the labeling algorithms. In fact, some of the best exact algorithms for the CVRP and VRPTW use SRCs [13,9,2,8]. A direct generalization of SRCs is to consider any Chvátal–Gomory rank 1 cut [6] obtained over a small subset of the rows. Preliminary experiments with those *Rank-1 Cuts (R1Cs)* were performed in [20].

The above mentioned algorithms separate SRCs in a very careful way, limiting the number of added cuts in order to avoid an excessive impact in the pricing. As a result, the potential of these cuts is not fully exploited. Pecin et al. [19] introduced the *limited memory* technique for making SRCs or R1Cs much less harmful to the labeling algorithms, without necessarily compromising their effectivity. The computational results on CVRP instances show a breakthrough: due to the improved bounds, the size of the largest solved literature instance increased from 150 to 360 customers. The improved bounds made possible by the limited memory also lead to big advances on VRPTW [18]. In this new context, there is a clear motivation for obtaining even better bounds by also separating more complex Rank-1 cuts.

This work is aimed at answering the following question: *what are the optimal sets of multipliers for Rank-1 Cuts with up to 5 rows?* This is done by a computational investigation of the set partitioning polyhedra. Experiments on CVRP instances show that the newly discovered cuts indeed lead to significantly improved bounds. Finally, we show how this allows to solve the Golden_20 instance (420 customers). The paper is organized as follows. Section 2 reviews SRCs, Rank-1 cuts and the limited memory technique. Section 3 describes the methodology used for discovering the optimal sets of multipliers. Section 4 presents experiments on CVRP instances. Finally, Section 5 contains some concluding remarks.

2. Limited memory Rank-1 Cuts

Given a base set $C \subseteq V$ and a multiplier p_i for each $i \in C$, the following valid inequality for SPF is a Rank-1 Cut (R1C):

$$\sum_{r\in\Omega}\left\lfloor\sum_{i\in C}p_i a_i^r\right\rfloor\lambda_r \le \left\lfloor\sum_{i\in C}p_i\right\rfloor. \tag{4}$$

The Subset Row Cuts (SRCs) were introduced in Jepsen et al. [13] and correspond to the particular case where all the multipliers have the same value $p = 1/k$, for some integer k. The following base set sizes and multipliers were investigated in that work:

- 3SRCs: $|C| = 3$ and $p = 1/2$. Those cuts have RHS 1 and can be viewed as weakened clique cuts. Nevertheless, they are still very effective in improving bounds and were the only SRCs actually separated in [13,9,2,8].
- 5,2SRCs: $|C| = 5$ and $p = 1/2$. Those cuts have RHS 2 and can be viewed as weakened odd hole cuts.

Pecin et al. [19] separated 3SRCs, 5,2SRCs and also:

- 4SRCs: $|C| = 4$ and $p = 2/3$. In spite of not having a multiplier of format $1/k$, they were still called SRCs in that work.
- 5,1SRCs: $|C| = 5$ and $p = 1/3$.

The only work that investigated and separated general R1Cs is Petersen et al. [20]. The used multiplier sets were of mod-k type [4], i.e, each individual multiplier should belong to $\{0, 1/k, \ldots, k - 1/k\}$ for some small integer k. Separation was performed exhaustively, testing all possibilities for certain values of C and k. The authors also tested finding the multipliers using the MIP formulation proposed in [10]. The experiments indicated that R1Cs more complex than SRCs could indeed improve bounds. However, those cuts could not be incorporated into state-of-the-art codes: not only the separation itself is very expensive, but the added cuts slow down too much the labeling algorithm, making the pricing intractable.

Given $C \subseteq V$, a vector of multipliers p of dimension $|C|$, a memory set M, $C \subseteq M \subseteq V$, the limited memory (C, M, p)-Rank 1 Cut (lm-R1C for short) is:

$$\sum_{r\in\Omega}\alpha(C, M, p, r)\lambda_r \le \left\lfloor\sum_{i\in C}p_i\right\rfloor, \tag{5}$$

where the coefficient of a route r is computed as:

function $\alpha(C, M, p, r)$
  coeff $\leftarrow 0$, *state* $\leftarrow 0$
  for every vertex $i \in r$ (in order) **do**
    if $i \notin M$ **then**
      state $\leftarrow 0$
    else if $i \in C$ **then**
      state $\leftarrow$ *state* $+ p_i$
      if *state* ≥ 1 **then**
        coeff $\leftarrow$ *coeff* $+ 1$, *state* $\leftarrow$ *state* $- 1$
return *coeff*

Variable *coeff* stores the coefficient to be returned. Each time a vertex i in C is visited, variable *state* is increased by p_i. When *state* becomes larger or equal to 1, its value is reduced by 1 unit and *coeff* is incremented. The previous definition is completely analogous to that of the limited memory Subset Row Cuts proposed (or lm-SRC) by Pecin et al. [19], the only difference being that the multiplier vector p replaces a single scalar, which was originally the same for all $i \in C$. When $M = V$, the Function α will return $\lfloor\sum_{i\in C} p_i a_i^r\rfloor$ and the lm-R1C will be identical to an R1C. On the other hand, when M is strictly contained in V, the lm-R1C may be a weakening of its corresponding R1C. This happens because every time the route r leaves M, the variable *state* is reset to zero, potentially decreasing the returned coefficient. Nevertheless, using a dynamic adjustment of the memory sets, the lm-R1Cs can still yield exactly the same gap improvements of ordinary R1Cs [19].

The advantage of the lm-R1Cs over R1Cs (or lm-SRCs over SRCs) is their much reduced impact on the labeling algorithm used in the pricing, when $|M| \ll |V|$. Labeling algorithms roughly consist of expanding sets of *non-dominated* partial routes called *buckets*, and then reducing the bucket sizes by eliminating newly dominated routes. Clearly, a crucial step of such algorithm is the bucket reduction by dominance, which heavily depends on the domination rule. Such a rule must ensure that the dominated partial route can always be replaced by the dominating one in any feasible route that contains it, without increasing (and potentially reducing) its cost. For the pricing subproblem that remains after adding lm-R1Cs to the SPF given by (1)–(3), the labeling algorithm keeps track of the state (the current value of variable *state* in the computation of α) of each lm-R1C with non-zero dual variable, for each partial route contained in each bucket. Then, for a partial route r_1 to dominate another partial route r_2, the reduced cost of r_2 must exceed that of r_1 at least in the amount of the sum of the dual variables of all lm-R1Cs whose states in r_1 are larger than those in r_2. This sum represents that worst-case increase in the reduced cost of r_1 that may not occur for r_2 considering the same feasible extension. If, on one hand, the limited memory technique greatly reduces the impact of lm-R1Cs on the pricing by resetting many states to zero, on the other hand, having a small number of possible states is still an essential property of these cuts to keep the pricing subproblem tractable. Note that this number is as small as the lowest common denominator of the components of the vector p for the newly proposed cuts.