Compact, Provably-Good LPs for Orienteering and Regret-Bounded Vehicle Routing*

Zachary Friggstad**[1] and Chaitanya Swamy***[2]

[1] Dept. of Computing Science, U. Alberta, Edmonton. zacharyf@ualberta.ca
[2] Combinatorics & Optimization, U. Waterloo, Waterloo. cswamy@uwaterloo.ca

Abstract. We develop polynomial-size LP-relaxations for *orienteering* and the *regret-bounded vehicle routing problem* (RVRP) and devise suitable LP-rounding algorithms that lead to various new insights and approximation results for these problems. In orienteering, the goal is to find a maximum-reward r-rooted path, possibly ending at a specified node, of length at most some given budget B. In RVRP, the goal is to find the minimum number of r-rooted paths of *regret* at most a given bound R that cover all nodes, where the regret of an r-v path is its length $- c_{rv}$. For *rooted orienteering*, we introduce a natural bidirected LP-relaxation and obtain a simple 3-approximation algorithm via LP-rounding. This is the *first LP-based* guarantee for this problem. We also show that *point-to-point* (P2P) *orienteering* can be reduced to a regret-version of rooted orienteering at the expense of a factor-2 loss in approximation. For RVRP, we propose two compact LPs that lead to significant improvements, in both approximation ratio and running time, over the approach in [10]. One is a natural modification of the LP for rooted orienteering; the other is an unconventional formulation motivated by certain structural properties of an RVRP-solution, which leads to a 15-approximation for RVRP.

1 Introduction

Vehicle-routing problems (VRPs) constitute a broad class of optimization problems that find a wide range of applications and have been widely studied in the Operations Research and Computer Science literature (see, e.g., [14, 18, 4, 2, 8]). Despite this extensive study, we have rather limited understanding of LP-relaxations for VRPs (with TSP and the minimum-latency problem, to a lesser extent, being exceptions), and this has been an impediment in the design of approximation algorithms for these problems.

Motivated by this gap in our understanding, we investigate whether one can develop polynomial-size (i.e., compact) LP-relaxations with good integrality gaps for VRPs, focusing on the fundamental *orienteering* problem [13, 4, 8] and the related *regret-bounded vehicle routing problem* (RVRP) [5, 10]. In *orienteering*, we are given rewards associated with clients located in a metric space, a length bound B, a start, and possibly end, location for the vehicle, and we

* A full version of the paper is available on the CS arXiv.

** Research supported by Canada Research Chairs program and NSERC Discovery Grant.

*** Research supported by NSERC grant 327620-09 and an NSERC DAS Award.

seek a route of length at most B that gathers maximum reward. This problem frequently arises as a subroutine when solving VRPs, both in approximation algorithms—e.g., for minimum-latency problems (MLPs) [3, 9, 6, 16], TSP with time windows [2], RVRP [5, 10]—as well as in computational methods where orienteering corresponds to the "pricing" problem encountered in solving set covering/partitioning LPs (a.k.a configuration LPs) for VRPs via a column-generation or branch-cut-and-price method. In RVRP, we have a metric space $\{c_{uv}\}$ on client locations, a start location r, and a *regret* bound R. The regret of a path P starting at r and ending at location v is $c(P) - c_{rv}$. The goal in RVRP is to find a minimum number of r-rooted paths of regret at most R that visit all clients.

Our contributions. We develop polynomial-size LP-relaxations for orienteering and RVRP and devise suitable rounding algorithms for these LPs, which lead to various new insights and approximation results for these problems.

In Section 3, we introduce a natural, compact LP-relaxation for *rooted orienteering*, wherein only the vehicle start node is specified, and design a simple rounding algorithm to convert an LP-solution to an integer solution losing a factor of at most 3 in the objective value. This is the *first LP-based* approximation guarantee for orienteering. In contrast, all other approaches for orienteering utilize dynamic programming (DP) to stitch together suitable subpaths.

In Section 4, we consider the more-general *point-to-point* (P2P) *orienteering* problem, where both the start and end nodes of the vehicle are specified. We present a novel reduction showing that P2P-orienteering can be *reduced* to a *regret-version* of rooted orienteering, wherein the length bound is replaced by a *regret bound*, incurring a factor-2 loss (Theorem 6). No such reduction to a rooted problem was known previously, and all known algorithms for P2P-orienteering rely on approximations to suitable P2P-path problems. Typically, constraining a VRP by requiring that routes include a fixed node t causes an increase in the route lengths of the unconstrained problem (as we need to attach t to the routes); this would violate the length bound in orienteering, but, notably, we devise a way to avoid this in our reduction. We believe that the insights gained from our reduction may find further application. Our results for rooted orienteering translate to the regret-version of orienteering, and combined with the above reduction, give a compact LP for P2P-orienteering having integrality gap at most 6.

Although we do not improve the current-best approximation factor of $(2 + \epsilon)$ for orienteering [8], we believe that our LP-based approach is nevertheless appealing for various reasons. First, our LP-rounding algorithms are quite simple, and arguably, simpler than the DP-based approaches in [4, 8]. Second, our LP-based approach offers the promising possibility that, by leveraging the key underlying ideas, one can obtain strong, compact LP-relaxations for other problems that utilize orienteering. Indeed, we already present evidence of such benefits by showing in Section 5.1 that our LP-insights for rooted orienteering yield a compact, provably-good LP for RVRP. (We remark that various configuration LPs considered for VRPs give rise to P2P-orienteering as the dual-separation problem, and utilizing our compact orienteering-LP in the dual could yield another way of obtaining a compact LP.) Finally, LP-based insights often tend to be powerful and have the potential to result in both improved guarantees, and algorithms for variants of the problem. In fact, we suspect that our orienteering LPs are better than what we have accounted for, and believe that they are a promising means of improving the state-of-the-art for orienteering.

Section 5 considers RVRP, and proposes two compact LP-relaxations for RVRP and corresponding rounding algorithms. Our LP-based algorithms not only yield improvements over the current-best 28.86-approximation for RVRP [10], but also result in substantial savings in running time compared to the algorithm in [10], which involves solving a configuration LP (with an exponential number of path variables) using the $\Omega(n^{1/\epsilon})$-time $(2+\epsilon)$-approximation algorithm for orienteering in [8] as a subroutine. The first LP for RVRP is a natural modification of our LP for rooted orienteering, which we show has integrality gap at most 27 (Theorem 7). In Section 5.2, we formulate a rather atypical LP-relaxation (R2) for RVRP by exploiting certain key structural insights for RVRP. We observe that an RVRP-solution can be regarded as a collection of distance-increasing rooted paths covering some *sentinel* nodes S and a low-cost way of connecting the remaining nodes to S, and our LP aims to find the best such solution. We design a rounding algorithm for this LP that leads to a 15-approximation algorithm for RVRP, which is a significant improvement over the guarantee obtained in [10].

Finally, in Section 6, we observe that our techniques imply that the integrality gap of a Held-Karp style LP for the *asymmetric-TSP* (ATSP) *path* problem is 2 for the class of asymmetric metrics induced by the regret objective.

To give an overview of our techniques, a key tool that we use in our rounding algorithms, which also motivates our LP-relaxations, is an arborescence-packing result of [1] showing that an r-preflow $x \in \mathbb{R}_+^A$ in a digraph $D = (N, A)$ (i.e., $x\big(\delta^{\text{in}}(v)\big) \geq x\big(\delta^{\text{out}}(v)\big) \ \forall v \neq r$) dominates a weighted collection of r-rooted (non-spanning) out-arborescences (Theorem 3). An r-preflow x in the bidirected version of our metric, D, is a natural relaxation of an r-rooted path, and the $r \rightsquigarrow u$ connectivity under x abstracts whether u lies on this path. This leads to our LP (R-O) for (rooted) orienteering. The idea behind the rounding is that if we know the node v on the optimum path with maximum c_{rv} value, then we can enforce that our the LP-preflow x is consistent with v. Hence, we can decompose x into arborescences containing v of average length at most B, which yield r-v paths of average *regret* at most $2(B - c_{rv})$. These in turn can be converted (see Lemma 1) into a weighted collection of paths of total weight at most 3, where each path has regret at most $B - c_{rv}$ and ends at some node u with $c_{ru} \leq c_{rv}$; returning the maximum-reward path in this collection yields a 3-approximation.

Related work. The orienteering problem seems to have been first defined in [13]. Blum et al. [4] gave the first $O(1)$-factor approximation for rooted orienteering. They obtained an approximation ratio of 4, which was generalized to P2P-orienteering, and improved to 3 [2] and then to $2 + \epsilon$ [8].

Orienteering is closely related to the k-{*stroll,* MST*,* TSP} problems, which seek a minimum-cost rooted {path,tree,tour} respectively spanning at least k nodes (so the roles of objective and constraint are interchanged). k-MST has a rich history of study that culminated in a factor-2 approximation for both k-MST and k-TSP [11]. Chaudhuri et al. [7] obtained a $(2+\epsilon)$-approximation algorithm for k-stroll. They also showed that for certain values of k, one can obtain a tree spanning k nodes and containing two specified nodes r, t, of cost at most the cheapest r-t path spanning k nodes. In particular, this holds for $k = n$, and yields an alternative way of obtaining a 2-approximation algorithm for the *minimum-regret TSP-path* problem considered in Section 6. The orienteering algorithms in [4, 2, 8] are all based on first obtaining suitable subpaths by approximating the *min-excess path* problem using a k-stroll algorithm as a subroutine, and then stitching together these subpaths via a DP. (For a rooted path, the notions of excess and regret coincide; we use the term regret as it is more in line with the terminology used in the vehicle-routing literature [17, 15].)

The use of regret as a vehicle-routing objective seems to have been first considered in [17], who present various heuristics, and RVRP is sometimes referred to as the *schoolbus problem* in the literature [17, 15, 5]. Bock et al. [5] were the first to consider RVRP from an approximation-algorithms perspective. They obtain approximation factors of $O(\log n)$ for general metrics and 3 for tree metrics. Subsequently, Friggstad and Swamy [10] gave the first constant-factor approximation algorithm for RVRP, obtaining a 28.86-approximation via an LP-rounding procedure for a configuration LP.

2 Preliminaries and notation

Both orienteering and RVRP involve a complete undirected graph $G = (\{r\} \cup V, E)$, where r is a distinguished root (or depot) node, and metric edge costs $\{c_{uv}\}$. Let $n = |V| + 1$. We call a path P in G rooted if it begins at r. We always think of the nodes on P as being ordered in increasing order of their distance along P from r, and directing P away from r means that we direct each edge $uv \in P$ from u to v if u precedes v (under this ordering). We use D_v to denote c_{rv} for all $v \in V \cup \{r\}$. Let $\mathcal{T}$ denote the collection of all r-rooted trees in G. For a vector $d \in \mathbb{R}^E$, and a subset $F \subseteq E$, we use $d(F)$ to denote $\sum_{e \in F} d_e$. Similarly, for a vector $d \in \mathbb{R}^V$ and $S \subseteq V$, we use $d(S)$ to denote $\sum_{v \in S} d_v$.

Regret metric and RVRP. For every ordered pair $u, v \in V \cup \{r\}$, define the *regret distance* (with respect to r) to be $c^{\mathsf{reg}}_{uv} := D_u + c_{uv} - D_v$. The regret distances $\{c^{\mathsf{reg}}_{uv}\}$ form an asymmetric metric that we call the *regret metric*. The regret of a node v lying on a rooted path P is given by $c^{\mathsf{reg}}_P(v) := c_P(v) - D_v = (c^{\mathsf{reg}}$-length of the r-v portion of $P)$, where $c_P(v)$ is the length of the r-v subpath of P. Define the regret of P to be $c^{\mathsf{reg}}(P)$, which is also the regret of the end-node of P. Observe that $c^{\mathsf{reg}}(Z) = c(Z)$ for any cycle Z. We utilize the following results from [10].

Lemma 1 ([10]). *Let $R \geq 0$. Given rooted paths $P_1, \ldots, P_k$ with total regret αR, we can efficiently find at most $k + \alpha$ rooted paths, each having regret at most R, that cover $\bigcup_{i=1}^k P_i$.*

Theorem 2 ([10]). *Let $x = (x_P)_{P \in \mathcal{P}}$ be a weighted collection of rooted paths such that $\sum_{P \in \mathcal{P}: v \in P} x_P \geq 1$ for all $v \in V$. Let $R \geq 0$ be some given parameter.*

Let $k = \sum_{P\in\mathcal{P}} x_P$ and $\sum_{P\in\mathcal{P}} c^{\mathsf{reg}}(P)x_P = \alpha R$. Then, for any $\theta \in (0,1)$, we can round x to obtain a collection of at most $\left(\frac{6}{1-\theta}+\frac{1}{\theta}\right)\alpha + \left\lceil\frac{k}{\theta}\right\rceil$ rooted paths each of regret at most R that cover all nodes in V.

Preflows and arborescence packing. Let $D = (\{r\}\cup V, A)$ be a digraph. We say that a vector $x \in \mathbb{R}_+^A$ is an *r-preflow* if $x\big(\delta^{\text{in}}(v)\big) \geq x\big(\delta^{\text{out}}(v)\big)$ for all $v \in V$. When r is clear from the context, we simply say preflow. A key tool that we exploit is an arborescence-packing result of Bang-Jensen et al. [1] showing that we can decompose a preflow into out-arborescences rooted at r, and this can be done in polytime [16]. By an *out-arborescence rooted at* r, we mean a subgraph B whose undirected version is a tree containing r, and where every node spanned by B except r has exactly one incoming arc in B.

Theorem 3 ([1, 16]). *Let $D = (\{r\}\cup V, A)$ be a digraph and $x \in \mathbb{R}_+^A$ be a preflow. Let $\lambda_v := \min_{\{v\}\subseteq S\subseteq V} x\big(\delta^{\text{in}}(S)\big)$ be the $r \rightsquigarrow v$ "connectivity" in D under capacities $\{x_a\}_{a\in A}$. Let $K > 0$ be rational. We can obtain out-arborescences $B_1,\dots,B_q$ rooted at r, and rational weights $\gamma_1,\dots,\gamma_q \geq 0$ such that $\sum_{i=1}^q \gamma_i = K$, $\sum_{i:a\in B_i}\gamma_i \leq x_a$ for all $a \in A$, and $\sum_{i:v\in B_i}\gamma_i = \min\big\{K,\lambda_v\big\}$ for all $v \in V$. Moreover, such a decomposition can be computed in time $\mathsf{poly}(|V|, \text{size of } K)$.*

3 Rooted orienteering

In the *rooted orienteering* problem, we have a complete undirected graph $G = (\{r\}\cup V, E)$, metric edge costs $\{c_{uv}\}$, a distance bound $B \geq 0$, and nonnegative node rewards $\{\rho(v)\}_{v\in V}$. The goal is to find a rooted path with cost at most B that collects the maximum reward. Whereas all current approaches for orienteering rely on a dynamic program to stitch together suitable subpaths, we present a simple LP-rounding-based 3-approximation algorithm for rooted orienteering.

Let $D = (\{r\}\cup V, A)$ denote the bidirected version of G, where both (u,v) and (v,u) get cost c_{uv}. To introduce our LP and our rounding algorithm, first suppose that we know a node v on the optimum path that has maximum distance D_v among all nodes on the optimum path. In our relaxation, we model the path as one unit of flow $x \in \mathbb{R}_+^A$ that exits r, visits only nodes u with $D_u \leq D_v$ and v to an extent of 1, and has cost at most B. Since we do not know the endpoint of our path, we relax x to be a preflow. Letting z_u^v denote the $r \rightsquigarrow u$ connectivity (under capacities $\{x_a\}$), the reward earned by x is $\mathsf{rewd}(x) := \sum_{u\in V}\rho(u)z_u^v$.

Our rounding procedure is based on the insight that Theorem 3 allows us to view x as a convex combination of arborescences, which we regard as r-rooted trees in G. Converting each tree into an r-v path (by standard doubling and shortcutting), we get a convex combination of rooted paths of average reward $\mathsf{rewd}(x)$, and average cost at most $2B - D_v$, and hence average c^{reg}-cost at most $2(B - D_v)$. Applying Lemma 1 to this collection, we then obtain a weighted collection of rooted paths of total weight at most 3 earning the same total reward, where each path has regret at most $B - D_v$, and hence, cost at most B (since it ends at some node u with $D_u \leq D_v$). Thus, the maximum-reward path in this collection yields a feasible solution with reward at least $\mathsf{rewd}(x)/3$.

Finally, we circumvent the need for "guessing" v by using variables z_v^v to indicate if v is the maximum-distance node on the optimum path. We impose that we have a preflow x^v of value z_v^v that visits v to an extent of z_v^v, and only visits nodes u with $D_u \leq D_v$, and z_u^v is now the $r \rightsquigarrow u$ connectivity under capacities x^v. (Note that $r \notin V$.)

$$\max \quad \sum_{u,v \in V} \rho(u) z_u^v \tag{R-O}$$

$$\text{s.t.} \quad x^v\big(\delta^{\text{in}}(u)\big) \geq x^v\big(\delta^{\text{out}}(u)\big) \qquad \forall u, v \in V \tag{1}$$

$$x^v\big(\delta^{\text{in}}(u)\big) = 0 \qquad \forall u, v \in V : D_u > D_v \tag{2}$$

$$x^v\big(\delta^{\text{in}}(S)\big) \geq z_u^v \qquad \forall v \in V, S \subseteq V, u \in S \tag{3}$$

$$\sum_{a \in A} c_a x_a^v \leq B z_v^v \qquad \forall v \in V \tag{4}$$

$$x^v\big(\delta^{\text{out}}(r)\big) = z_v^v \quad \forall v \in V, \qquad \sum_v z_v^v = 1, \quad x, z \geq 0.$$

This formulation can be converted to a compact LP by introducing flow variables $f^{u,v} = \{f_a^{u,v}\}_{a \in A}$, and encoding the cut constraints (3) by imposing that $f^{u,v} \leq x^v$, and that $f^{u,v}$ sends z_u^v units of flow from r to u. Observe that: (a) if $D_u > D_v$ then $z_u^v \leq x^v\big(\delta^{\text{in}}(u)\big) = 0$; (b) we have $z_u^v \leq x^v\big(\delta^{\text{in}}(V)\big) = x^v\big(\delta^{\text{out}}(r)\big) = z_v^v$ for all u, v. Let (x^*, z^*) be an optimal solution to (R-O), of value OPT.

Theorem 4. *We can round (x^*, z^*) to a rooted-orienteering solution of value at least $OPT/3$.*

Proof. For each v with $z_v^{*v} > 0$ we apply Theorem 3 with $K = z_v^{*v}$ to obtain r-rooted out-arborescences, which we view as rooted trees in G, and associated nonnegative weights $\{\gamma_T^v\}_{T \in \mathcal{T}}$; recall that $\mathcal{T}$ is the collection of all r-rooted trees. So we have $\sum_T \gamma_T^v = z_v^{*v}$, $\sum_T \gamma_T^v c(T) \leq \sum_a c_a x_a^{*v} \leq B z_v^{*v}$, and $\sum_{T: u \in T} \gamma_T^v \geq z_u^{*v}$ for all $u \in V$. Note that for every T with $\gamma_v^T > 0$, we have $v \in T$, and $D_u \leq D_v$ for all $u \in T$ (as otherwise, we have $x^{*_v}\big(\delta^{\text{in}}(u)\big) = 0$). For every v and every tree T with $\gamma_T^v > 0$, we do the following. First, we double the edges not lying on the r-v path of T and shortcut to obtain a simple r-v path P_T^v. So

$$\sum_T \gamma_T^v c^{\mathsf{reg}}(P_T^v) \leq 2 \sum_T \gamma_T^v \big(c(T) - D_v\big) = 2 z_v^{*v}(B - D_v). \tag{5}$$

Next, we use Lemma 1 with regret-bound $B - D_v$ to break P_T^v into a collection $\mathcal{P}_T^v$ of at most $1 + \frac{c^{\mathsf{reg}}(P_T^v)}{B - D_v}$ rooted paths, each having c^{reg}-cost at most $B - D_v$. Note that if $B = D_v$, then $c^{\mathsf{reg}}(P_T^v) = 0$, and we use the convention that $0/0 = 0$, so $|\mathcal{P}_T^v| = 1$ in this case. Each path in $\mathcal{P}_T^v$ ends at a vertex u with $D_u \leq D_v$, so its c-cost is at most B. Now, for all $v \in V$, we have

$$\sum_T \gamma_T^v \sum_{P \in \mathcal{P}_T^v} \rho(P) = \sum_T \gamma_T^v \rho(P_T^v) \geq \sum_u \rho(u) z_u^{*v} \tag{6}$$

$$\sum_T \gamma_T^v |\mathcal{P}_T^v| \leq \sum_T \gamma_T^v \Big(1 + \tfrac{c^{\mathsf{reg}}(P_T^v)}{B - D_v}\Big) \leq z_v^{*v} + 2 z_v^{*v} = 3 z_v^{*v} \tag{7}$$

where the last inequality in (7) follows from (5). Therefore, the maximum-reward path in $\bigcup_{v,T:\gamma_T^v>0}\mathcal{P}_T^v$ earns reward at least

$$\Big(\sum_{v,T}\gamma_T^v\sum_{P\in\mathcal{P}_T^v}\rho(P)\Big)\Big/\Big(\sum_{v,T}\gamma_T^v|\mathcal{P}_T^v|\Big)\geq\frac{\sum_{v,u}\rho_u z_u^{*v}}{3\sum_v z_v^{*v}}=OPT/3.\qquad\square$$

The following variant of rooted orienteering, which we call *regret orienteering*, will be useful in Section 4. In regret orienteering, instead of a cost bound B, we are given a *regret bound* R, and we seek a rooted path of *regret* at most R that collects the maximum reward. The LP-relaxation for regret-orienteering is very similar to (R-O); the only changes are that z_v^v now indicates if v is the *end node* of the optimum path, and so we drop (2) and replace (4) with $\sum_{a\in A}c_a x_a^v\leq(D_v+R)z_v^v$. The rounding algorithm is essentially unchanged: we convert the trees obtained from x^v into r-v paths, which are then split into paths of regret at most R. Theorem 4 yields the following corollary.

Corollary 5. *There is an LP-based* 3*-approximation for regret orienteering.*

4 Point-to-point orienteering

We now consider the generalization of rooted orienteering, where we have a start node r *and* an end node t, and we seek an r-t path with cost at most B that collects the maximum reward. We may assume that r and t have 0 reward, i.e., $\rho(r)=\rho(t)=0$. The main result of this section is a novel reduction showing that *point-to-point (*P2P*) orienteering* problem can be *reduced* to regret orienteering losing a factor of at most 2 (Theorem 6). Combining this with our LP-approach for regret orienteering and Corollary 5, we obtain an LP-relaxation for P2P-orienteering having integrality gap at most 6 (described in the full version). We believe that the insights gained from this reduction may find further application.

Theorem 6. *An α-approximation algorithm for regret orienteering (where $\alpha\geq 1$) can be used to obtain a 2α-approximation algorithm for* P2P*-orienteering.*

Proof. Let $\big(G=(\{r,t\}\cup V,E),\{c_{uv}\},\{\rho(u)\},B\big)$ be a P2P-orienteering instance. Our reduction is simple. Let P^* be an optimal solution. We "guess" a node $v\in P^*$ (which could be r or t) such that $D_v+c_{vt}=\max_{u\in P^*}(D_u+c_{ut})$. (That is, we enumerate over all choices for v.) Let $S=\{u\in\{r,t\}\cup V: D_u+c_{ut}\leq D_v+c_{vt}\}$. We then consider two regret orienteering problems, both of which have regret bound $R=B-D_v-c_{vt}$ and involve only nodes in S (i.e., we equivalently set $\rho(u)=0$ for all $u\notin S$); the first problem has root r, and the second has root t. Let P_1 and P_2 be the solutions obtained for these two problems respectively by our α-approximation algorithm. So for some $u_1,u_2\in S$, P_1 is an r-u_1 path, and P_2 may be viewed as a u_2-t path. Notice that P_1 appended with the edge u_1t yields an r-t path of *cost* at most $D_{u_1}+c^{\mathsf{reg}}(P_1)+c_{u_1t}\leq D_{u_1}+c_{u_1t}+B-D_v-c_{vt}\leq B$, since $u_1\in S$. Similarly P_2 appended with the edge ru_2 yields an r-t path of cost at most B. We return P_1+u_1t or ru_2+P_2, whichever has higher reward.

To analyze this, we observe that the r-v portion of P^* is a feasible solution to the regret-orienteering instance with root r, since its cost is at most $B - c_{vt}$, and hence, its regret is at most R. Similarly, the v-t portion of P^* (viewed in reverse) is a feasible solution to the regret-orienteering instance with root t. Therefore, $\max\{\rho(P_1 + u_1 t), \rho(r u_2 + P_2)\} \geq \rho(P^*)/2\alpha$. □

5 Compact LPs and improved guarantees for RVRP

Recall that in the *regret-bounded vehicle routing problem* (RVRP), we are given an undirected complete graph $G = (\{r\} \cup V, E)$ on n nodes with a distinguished root (depot) node r, metric edge costs or distances $\{c_{uv}\}$, and a regret-bound R. The goal is to find the minimum number of rooted paths that cover all nodes so that the regret of each node with respect to the path covering it is at most R. Throughout, let O^* denote the optimal value of the RVRP instance. We describe two compact LP-relaxations for RVRP and corresponding rounding algorithms that yield improvements, in both approximation ratio and running time, over the RVRP-algorithm in [10]. In Section 5.1, we observe that the compact LP for orienteering (R-O) yields a natural LP for RVRP; by combining the rounding ideas used for orienteering and Theorem 2, we obtain a 27-approximation algorithm for RVRP. In Section 5.2, we formulate an unorthodox, stronger LP-relaxation (R2) for RVRP by leveraging some key structural insights in [10]. We devise a rounding algorithm for this LP that leads to a 15-approximation algorithm for RVRP, which is a significant improvement over the guarantee obtained in [10].

5.1 Extending the orienteering LP to RVRP

The LP-relaxation below can be viewed as a natural variant of the orienteering LP adapted to RVRP. As before, let $D = (\{r\} \cup V, A)$ be the bidirected version of G. For each node v, x^v is a preflow (constraint (8)) of value z^v_v such that the $r \rightsquigarrow u$ connectivity under capacities $\{x^v_a\}$ is at least z^v_u for all u, v (constraint (9)). As before, we can obtain a compact formulation by replacing the cut constraints (9) with constraints involving suitable flow variables.

$$\min \quad \sum_v z^v_v \tag{R1}$$

$$\text{s.t.} \quad x^v\big(\delta^{\text{in}}(u)\big) \geq x^v\big(\delta^{\text{out}}(u)\big) \qquad \forall u, v \in V \tag{8}$$

$$x^v\big(\delta^{\text{in}}(S)\big) \geq z^v_u \qquad \forall v \in V, S \subseteq V, u \in S \tag{9}$$

$$\sum_{a \in A} c_a x^v_a \leq (D_v + R) z^v_v \qquad \forall v \in V$$

$$x^v\big(\delta^{\text{out}}(r)\big) = z^v_v \quad \forall v \in V, \qquad \sum_{v \in V} z^v_u \geq 1 \quad \forall u \in V, \qquad x, z \geq 0.$$

Theorem 7. *We can round an optimal solution to* (R1) *to obtain a* 27*-approximation for* RVRP*.*

5.2 A new compact LP for RVRP leading to a 15-approximation

We now propose a different LP for RVRP, which leads to a much-improved 15-approximation for RVRP. To motivate this LP, we first collect some facts from [10, 4] pertaining to the regret objective. By merging all nodes at distance 0 from each other, we may assume that $c_{uv} > 0$ for all $u, v \in V \cup \{r\}$, and hence $D_v > 0$ for all $v \in V$.

Definition 8 ([10]). Let P be a rooted path ending at w. Consider an edge (u, v) of P, where u precedes v on P. We call this a *red* edge of P if there exist nodes x and y on the r-u portion and v-w portion of P respectively such that $D_x \geq D_y$; otherwise, we call this a *blue* edge of P. (Note that the first edge of P is always a blue edge.)

For a node $x \in P$, let $\mathsf{red}(x, P)$ denote the maximal subpath Q of P containing x consisting of only red edges (which might be the trivial path $\{x\}$). Call the collection $\{\mathsf{red}(x, P) : x \in P\}$ of subpaths, the red intervals of P.

Lemma 9 ([4]). *For any rooted path P, we have $\sum_{e \text{ red on } P} c_e \leq \frac{3}{2} c^{\mathsf{reg}}(P)$.*

Lemma 10 ([10]). *(i) Let u, v be nodes on a rooted path P such that u precedes v on P and $\mathsf{red}(u, P) \neq \mathsf{red}(v, P)$; then $D_u < D_v$. (ii) If P' is obtained by shortcutting P so that it contains at most one node from each red interval of P, then for every edge (x, y) of P' with x preceding y on P', we have $D_x < D_y$.*

We say that a node u on a rooted path of P is a *sentinel* of P if u is the first node of $\mathsf{red}(u, P)$. Part (ii) above shows that if we shortcut each path P of an optimal RVRP-solution past the non-sentinel nodes of P, then we obtain a distance-increasing collection of paths. Moreover, part (i) implies that if x and y are sentinels on P with x appearing before y, then $\max_{u \in \mathsf{red}(x,P)} D_u < \min_{u \in \mathsf{red}(y,P)} D_u$. Finally, every non-sentinel node is connected to the sentinel corresponding to its red interval via red edges, and Lemma 9 shows that the total (c-) cost of these edges at most $1.5R$(optimal value).

Thus, we can view an RVRP-solution as a collection of distance-increasing rooted paths covering some sentinel nodes S, and a low-cost way of connecting the nodes in $V \setminus S$ to S. Our LP-relaxation searches for the best such solution. Let $\mathcal{D} := \{D_v : v \in V\}$. For every $u \in V$, define $\mathcal{D}_u$ to be the collection $\{[d_1, d_2] : d_1, d_2 \in \mathcal{D},\ d_1 \leq D_u \leq d_2\}$ of (closed) intervals. We have variables $x_{u,I,u}$ for every node $u \in V$ and interval $I = [d_1, d_2] \in \mathcal{D}_u$ to indicate if u is a sentinel and d_1, d_2 are the minimum and maximum distances (from r) respectively of nodes in the red interval corresponding to u; we say that I is u's distance interval. We also have variables $x_{u,I,v}$ for $v \neq u$ to indicate that v is connected to sentinel u with distance interval I, and edge variables $\{z_e\}_{e \in E}$ that encode these connections. Finally, we have flow variables $f_{r,u,I}, f_{u,I,v,J}, f_{u,I,t}$ for all $u, v \in V$ and $I \in \mathcal{D}_u$, $J \in \mathcal{D}_v$ that encode the distance-increasing rooted paths on the sentinels, with t representing a fictitious sink. We include constraints that encode that the distance intervals of sentinels lying on the same path are disjoint, and a non-sentinel v can be connected to (u, I) only if $D_v \in I$. We obtain the following LP.

$$\min \quad \sum_{u\in V, I\in\mathcal{D}_u} f_{r,u,I} \tag{R2}$$

$$\text{s.t.} \quad \sum_{u\in V, I\in\mathcal{D}_u} x_{u,I,v} \geq 1 \qquad \forall v \in V \tag{10}$$

$$x_{u,I,v} \leq x_{u,I,u}, \quad x_{u,I,v} = 0 \quad \text{if } D_v \notin I \qquad \forall u,v \in V, I \in \mathcal{D}_u \tag{11}$$

$$z\big(\delta(S)\big) \geq \sum_{u\notin S, I\in\mathcal{D}_u} x_{u,I,v} \qquad \forall v \in V, \{v\} \subseteq S \subseteq V \tag{12}$$

$$f_{r,u,I} + \sum_{v\in V, J\in\mathcal{D}_v} f_{v,J,u,I} = x_{u,I,u} \qquad \forall u \in V, I \in \mathcal{D}_u \tag{13}$$

$$\sum_{v\in V, J\in\mathcal{D}_v} f_{u,I,v,J} + f_{u,I,t} = x_{u,I,u} \qquad \forall u \in V, I \in \mathcal{D}_u \tag{14}$$

$$f_{u,I,v,J} = 0 \qquad \forall u,v \in V, I \in \mathcal{D}_u, J \in \mathcal{D}_v : I \cap J \neq \emptyset \text{ or } D_v \leq D_u \tag{15}$$

$$\sum_{u,v\in V, I\in\mathcal{D}_u, J\in\mathcal{D}_v} c^{\mathsf{reg}}_{uv} f_{u,I,v,J} \leq R \cdot \sum_{u\in V, I\in\mathcal{D}_u} f_{r,u,I} \tag{16}$$

$$\sum_{e\in E} c_e z_e \leq 1.5R \cdot \sum_{u\in V, I\in\mathcal{D}_u} f_{r,u,I} \tag{17}$$

$$x, z, f \geq 0.$$

Constraint (10) encodes that every node v is either a sentinel or is connected to a sentinel; (11) ensures that if v is assigned to (u, I), then u is indeed a sentinel with distance interval I and that $D_v \in I$. Constraints (12) ensure that the z_es (fractionally) connect each non-sentinel v to the sentinel specified by the $x_{u,I,v}$ variables. Constraints (13), (14) encode that each sentinel (u, I) lies on rooted paths, and (15) ensures that these paths are distance increasing and moreover the distance intervals of the sentinels on the paths are disjoint. Finally, letting k denote the number of paths used, (16), (17) encode that the total regret of the distance-increasing paths is at most kR (note that $c^{\mathsf{reg}}_{ru} = 0$ for all u), and the total cost of the edges used to connect non-sentinels to sentinels is at most $1.5kR$. As before, the cut constraints (12) can be equivalently stated using flows to obtain a polynomial-size LP. Let (x^*, z^*, f^*) be an optimal solution to (R2) and OPT denote its objective value. We have already argued that an optimal RVRP-solution yields an integer solution to (R2), so we obtain that $\lceil OPT \rceil$ is at most the optimal value, O^*, of the RVRP instance.

Our rounding algorithm proceeds in a similar fashion as the RVRP-algorithm in [10]; yet, we obtain an improved guarantee since one can solve (R2) exactly whereas one can only obtain a $(2+\epsilon)$-approximate solution to the configuration LP in [10]. Let $\theta \in (0, 1)$ be a parameter that we will set later. We first obtain a forest F of c-cost at most $\frac{3R}{1-\theta} \cdot OPT$ such that every component Z contains a witness node v that is assigned to an extent of at least θ to sentinels in Z. We argue that if we contract the components of F, then the distance-increasing sentinel flow paths yield an acyclic flow that covers every contracted component to an extent of at least θ. Hence, using the integrality property of flows, we obtain an integral flow, and hence a collection of at most $\lceil \frac{OPT}{\theta} \rceil$ rooted paths, that covers every component and has cost at most $\frac{R}{\theta} \cdot OPT$. Next, we show that we can uncontract the components and attach the component-nodes to these rooted paths incurring an additional cost of at most $\frac{6R}{1-\theta} \cdot OPT$. Finally, by applying Lemma 1, we obtain an RVRP solution with at most $\left(\frac{6}{1-\theta} + \frac{1}{\theta}\right) OPT + \lceil \frac{OPT}{\theta} \rceil$ rooted paths.

A1. For $S \subseteq V$, define $h(S) = 1$ if $\sum_{u \in S, I \in \mathcal{D}_u} x^*_{u,I,v} < \theta$ for all $v \in S$, and 0 otherwise. h is a *downwards-monotone* cut-requirement function: if $\emptyset \neq A \subseteq B$, then $h(A) \geq h(B)$. Use the LP-relative 2-approximation algorithm in [12] for $\{0,1\}$ downwards-monotone functions to obtain a forest F such that $|\delta(S) \cap F| \geq h(S)$ for all $S \subseteq V$.

A2. For every component Z of F with $r \notin Z$, pick a *witness node* $w \in Z$ such that $\sum_{u \in Z, I \in \mathcal{D}_u} x^*_{u,I,w} \geq \theta$. Let $\sigma(w) = \{(u,I) : u \in Z, x^*_{u,I,w} > 0\}$. Let $W \subseteq V$ be the set of all such witness nodes.

A3. f^* is an $r \rightsquigarrow t$ flow in an auxiliary graph having nodes r, t, and (u,I) for all $u \in V, I \in \mathcal{D}_u$, edges $(r,(u,I))$, $((u,I),t)$ for all $u \in V, I \in \mathcal{D}_u$, and edges $((u,I),(v,J))$ for all $u,v \in V, I \in \mathcal{D}_u, J \in \mathcal{D}_v$ such that $D_u < D_v$ and $I \cap J = \emptyset$. Let $\{f^*_P\}_{P \in \mathcal{P}}$ be a path-decomposition of this flow. Modify each flow path $P \in \mathcal{P}$ as follows. First, drop t from P. Shortcut P past the nodes in P that are not in $\{r\} \cup \bigcup_{w \in W} \sigma(w)$. The resulting path maps naturally to a rooted path in G (obtained by simply dropping the distance intervals), which we denote by $\pi(P)$. Clearly, $c^{\mathsf{reg}}\big(\pi(P)\big) \leq \sum_{((u,I),(v,J)) \in P} c^{\mathsf{reg}}_{uv}$ since shortcutting does not increase the regret cost.

A4. Let $\mathcal{Q}$ be the collection of rooted paths obtained by taking the paths $\{\pi(P) : P \in \mathcal{P}\}$ and contracting the components of F. Let H be the directed graph (which we prove is acyclic) obtained by directing the paths in $\mathcal{Q}$ away from r. To avoid notational clutter, for a component Z of F, we use Z to also denote the corresponding contracted node in H. For each $Q \in \mathcal{Q}$, define $y_Q = \sum_{P \in \mathcal{P}: \pi(P) \text{ maps to } Q} f^*_P$.

A5. Use the integrality property of flows to round the flow $\left\{\frac{y_Q}{\theta}\right\}_{Q \in \mathcal{Q}}$ to an integer flow of value $k \leq \lceil \frac{OPT}{\theta} \rceil$ and regret-cost at most $\frac{R}{\theta} \cdot OPT$. Since H is acyclic, this yields rooted paths $\hat{P}_1, \ldots, \hat{P}_k$ so that every component Z of F lies on exactly one $\hat{P}_i$ path.

A6. We map the $\hat{P}_i$s to rooted paths in G that cover V as follows. Consider a path $\hat{P}_i$. Let Z be a component lying on $\hat{P}_i$, and $u,v \in Z$ be the nodes where $\hat{P}_i$ enters and leaves Z respectively. We add to $\hat{P}_i$ a u-v path that covers all nodes of Z obtained by doubling all edges of Z except those on the u-v path in Z and shortcutting. Let $\tilde{P}_i$ be the rooted path in G obtained by doing this for all components lying on $\hat{P}_i$.

A7. Finally, we use Lemma 1 to convert $\tilde{P}_1, \ldots, \tilde{P}_k$ to an RVRP-solution.

Theorem 11. *The above algorithm returns an* RVRP*-solution with at most* $\left(\frac{6}{1-\theta} + \frac{1}{\theta}\right) OPT + \lceil \frac{OPT}{\theta} \rceil$ *paths. Thus, taking* $\theta = \frac{1}{3}$*, we obtain at most* $15 \cdot O^*$ *paths.*

6 Minimum-regret TSP-path

We now consider the *minimum-regret TSP-path* problem, wherein we have (as before), a complete graph $G = (V', E)$, $r, t \in V'$, metric edge costs $\{c_{uv}\}$, and we seek a minimum-regret r-t path that visits all nodes. Observe that this is precisely the ATSP-*path* problem under the asymmetric regret metric c^{reg}. We establish a tight bound of 2 on the integrality gap of the standard ATSP-path LP for the class of regret-metrics (induced by a symmetric metric). We consider the following LP for min-regret TSP path. Let $D = (V', A)$ be the bidirected version of G. Let $b_t = 1 = -b_r$ and $b_v = 0$ for all $v \in V' \setminus \{r, t\}$.

$$\min \sum_{a \in A} c_a^{\mathsf{reg}} x_a \quad \text{s.t.} \quad x\big(\delta^{\text{in}}(v)\big) - x\big(\delta^{\text{out}}(v)\big) = b_v \quad \forall v \in V', \quad x \geq 0 \quad \text{(R-TSP)}$$
$$x(\delta^{\text{in}}(S)) \geq 1 \quad \forall \emptyset \neq S \subseteq V \setminus \{r\}.$$

Theorem 12. *The integrality gap of* (R-TSP) *is 2 for regret metrics, and we can obtain in polytime a Hamiltonian r-t path P with $c^{\mathsf{reg}}(P) \leq 2 \cdot OPT_{R\text{-}TSP}$.*

References

1. J. Bang-Jensen, A. Frank, and B. Jackson. Preserving and increasing local edge-connectivity in mixed graphs. *SIAM J. Discrete Math.*, 8(2):155–178, 1995.
2. N. Bansal, A. Blum, S. Chawla, and A. Meyerson. Approximation algorithms for deadline-TSP and vehicle routing with time windows. In *36th STOC*, 2004.
3. A. Blum, P. Chalasani, D. Coppersmith, B. Pulleyblank, P. Raghavan, and M. Sudan. The Minimum Latency Problem. In *26th STOC*, pages 163–171, 1994.
4. A. Blum, S. Chawla, D. R. Karger, T. Lane, and A. Meyerson. Approximation algorithms for orienteering and discount-reward TSP. *SICOMP*, 37:653–670, 2007.
5. A. Bock, E. Grant, J. Könemann, and L. Sanita. The school bus problem in trees. In *Proceedings of ISAAC*, 2011.
6. D. Chakrabarty and C. Swamy. Facility location with client latencies: linear-programming based techniques for minimum-latency problems. *Mathematics of Operations Research*, 41(3):865–883, 2016.
7. K. Chaudhuri, P. B. Godfrey, S. Rao, and K. Talwar. Paths, Trees and Minimum Latency Tours. In *Proceedings of 44th FOCS*, pages 36–45, 2003.
8. C. Chekuri, N. Korula, and M. Pál. Improved algorithms for orienteering and related problems. *ACM Transactions on Algorithms*, 8(3), 2012.
9. J. Fakcharoenphol, C. Harrelson, and S. Rao The k-traveling repairman problem. *ACM Trans. on Alg.*, Vol 3, Issue 4, Article 40, 2007.
10. Z. Friggstad and C. Swamy. Approximation algorithms for regret-bounded vehicle routing and applications to distance-constrained vehicle routing. In *Proceedings of STOC*, pages 744–753, 2014. Detailed version posted on *CS arXiv*, Nov 2013.
11. N. Garg. Saving an epsilon: a 2-approximation for the k-MST problem in graphs. In *Proceedings of the 37th STOC*, pages 396–402, 2005.
12. M. X. Goemans and D. P. Williamson Approximating minimum-cost graph problems with spanning tree edges. *Operations Research Letters* 16:183–189, 1994.
13. B. L. Golden, L. Levy, and R. Vohra. The orienteering problem. *Naval Research Logistics*, 34:307–318, 1987.
14. M. Haimovich and A. Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10:527–542, 1985.
15. J. Park and B. Kim. The school bus routing problem: A review. *European Journal of Operational Research*, 202(2):311–319, 2010.
16. I. Post and C. Swamy. Linear-programming based techniques for multi-vehicle minimum latency problems. In *Proceedings of 26th SODA*, pages 512–531, 2015.
17. M. Spada, M. Bierlaire, and T. Liebling. Decision-aiding methodology for the school bus routing and scheduling problem. *Transportation Sc.*, 39:477–490, 2005.
18. P. Toth and D. Vigo, eds. *The Vehicle Routing Problem.* SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2002.