OPERATIONS RESEARCH
Vol. 60, No. 1, January–February 2012, pp. 123–127
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.1110.0967

Technical Note

Approximation Algorithms for VRP with Stochastic Demands

Anupam Gupta
Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, anupamg@cs.cmu.edu

Viswanath Nagarajan
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, viswanath@us.ibm.com

R. Ravi
Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, ravi@cmu.edu

We consider the *vehicle routing problem with stochastic demands* (VRPSD). We give randomized approximation algorithms achieving approximation guarantees of $1+\alpha$ for *split-delivery* VRPSD, and $2+\alpha$ for *unsplit-delivery* VRPSD; here α is the best approximation guarantee for the traveling salesman problem. These bounds match the best known for even the respective deterministic problems [Altinkemer, K., B. Gavish. 1987. Heuristics for unequal weight delivery problems with a fixed error guarantee. *Oper. Res. Lett.* **6**(4) 149–158; Altinkemer, K., B. Gavish. 1990. Heuristics for delivery problems with constant error guarantees. *Transportation Res.* **24**(4) 294–297]. We also show that the "cyclic heuristic" for split-delivery VRPSD achieves a constant approximation ratio, as conjectured in Bertsimas [Bertsimas, D. J. 1992. A vehicle routing problem with stochastic demand. *Oper. Res.* **40**(3) 574–585].

Subject classifications: analysis of algorithms: suboptimal algorithms; transportation: vehicle routing; networks/graphs: stochastic.
Area of review: Optimization.
History: Received June 2010; revision received November 2010; accepted January 2011.

1. Introduction

The capacitated vehicle routing problem (VRP) is defined on a finite metric space (V, d), where V is a finite set of locations/vertices and d: $V \times V \to \mathbb{R}_+$ a distance function that is symmetric and satisfies the triangle inequality. There is a specified *depot* location $r \in V$, and the problem involves distributing (identical) items from the depot to other locations. Specifically, the depot r has an infinite supply of items, and a single vehicle of capacity $Q \geqslant 0$ (initially located at the depot r) is used to distribute the items. There are $q_i \in \{0, 1, \ldots, Q\}$ units of the item demanded at each location $i \in V$. The objective is to find a minimum-length tour of the vehicle that satisfies demands at all locations subject to the constraint that the vehicle carries at most Q units at any time. In the *split-delivery* version of the problem, the demand at a location may be satisfied by multiple visits of the vehicle. In the *unsplit-delivery* version, the entire demand at a location must be satisfied in a single visit by the vehicle.

The vehicle routing problem with stochastic demands (VRPSD) involves demands that are random variables with known distributions: in particular, the demand at each location $i \in V$ is specified by random variable ξ_i (assumed to be in the range $[0, Q]$). The random variables $\{\xi_i\}_{i\in V}$ are independent of each other and of the vehicle routing algorithm. The exact demand at any location i (i.e., the instantiation of the random variable ξ_i) is known only when the vehicle visits its location i. A feasible policy for the vehicle is any strategy of visiting locations such that all demands are satisfied. The goal is to design a policy that minimizes the *expected* length of the vehicle route. Again there are two variants of the problem, allowing for split and unsplit deliveries; we consider both of them. VRPSD has been studied extensively in the literature, and several heuristics have been proposed. For surveys on VRPSD, see, e.g., Stewart and Golden (1983), Powell et al. (1995), Gendreau et al. (1996), Bertsimas and Simchi-Levi (1996), and Dror (2002).

Related Work. Throughout this paper, we let $\alpha \geqslant 1$ denote the best approximation guarantee for the *Traveling Salesman Problem*. We have $\alpha = \frac{3}{2}$ for general metrics (Christofides 1977), and $\alpha = 1 + \epsilon$ (for any constant $\epsilon > 0$) for constant dimensional Euclidean metrics (Arora 1998, Mitchell 1999). The best known approximation guarantees for split-delivery VRP is $1 + \alpha \cdot (1 - 1/Q)$ (Haimovich and Rinnooy Kan 1985, Altinkemer and Gavish 1990), and for unsplit-delivery VRP is $2 + \alpha \cdot$

$(1-2/Q)$ (Altinkemer and Gavish 1987); recall that Q is an upper bound on the demand at any vertex. These bounds have been improved slightly to $1+\alpha\cdot(1-1/Q)-1/(3Q^3)$ and $2+\alpha\cdot(1-2/Q)-1/(3Q^3)$, respectively, (when $Q\geqslant 3$) (Bompadre et al. 2006). Recently, some improved bounds have been obtained for unit-demand VRP on the Euclidean plane: Bompadre et al. (2007) showed that the algorithm of Haimovich and Rinnooy Kan (1985) achieves an approximation ratio $2-c$ for random instances (here $c>0$ is an absolute constant). Also, Das and Mathieu (2010) gave a quasipolynomial time approximation scheme (i.e., an algorithm, that for any constant $\epsilon>0$, achieves an $(1+\epsilon)$-approximation with running time $2^{(\log n)^{O(1/\epsilon)}}$) for worst-case instances of points in the Euclidean plane.

Moving to stochastic demands, a $(1+\alpha+o(1))$-approximation algorithm is known for split-delivery VRPSD in the special case of *identical demand distributions* (Bertsimas 1992). For the case of general distributions, the algorithm in Bertsimas (1992) was shown to be a $Q+\alpha$ approximation, and obtaining tighter bounds was left open.

Our Results. In this note, we present a simple randomized approximation algorithm for VRPSD achieving the following worst-case guarantees: $(1+\alpha)$ for *split-delivery VRPSD* (Theorem 3.1), and $(2+\alpha)$ for *unsplit-delivery VRPSD* (Theorem 4.1). This matches (up to an additive $O(\alpha/Q)$ term) the corresponding best known guarantees for deterministic VRP. We also show in Theorem 3.2 that the "cyclic heuristic" for VRPSD suggested in Bertsimas (1992) achieves a constant approximation guarantee, as conjectured.

We note that our algorithms for VRPSD yield *a priori* strategies, that involve visiting vertices in some fixed order (in fact, according to any α-approximate TSP tour). Furthermore, the algorithms do not even require knowledge of the demand distributions at different vertices: it suffices to know just which vertices have a demand distribution that is not identically zero. The idea of visiting vertices in the order given by an approximate TSP tour is a natural strategy (this was also used in Bertsimas 1992). The key idea in our analysis is to condition on the demand realization—this allows us to bound the total length of refill-trips in terms of known lower bounds for the problem. Furthermore, we use randomization to obtain tighter approximation guarantees: the randomness permits an averaging argument similar to Haimovich and Rinnooy Kan (1985), Altinkemer and Gavish (1990, 1987) in spite of having stochastic demands. Due to the randomization, our solutions have the somewhat counterintuitive property that while refilling items at the depot, they fill up to a level *strictly below* the maximum capacity Q. We note that for both split and unsplit deliveries, the natural strategy of always refilling to maximum capacity Q also achieves a constant approximation guarantee. However, the constants are weaker than those for the randomized strategies (compare Theorem 3.2 to Theorem 3.1).

2. Preliminaries

The following discussion holds for both split and unsplit delivery versions of VRPSD.

Without loss of generality, we may assume that none of the demand random variables is identically zero, because such locations can just be ignored. Under this assumption, any feasible policy must visit every location with probability 1; otherwise, because demands are independent, there is a nonzero probability that some demand is not satisfied, implying that the policy is infeasible. Consequently, the minimum-length traveling salesman tour on all locations in metric (V,d) is a lower bound for VRPSD.

Conditioned on the demand realization q_i at each location $i\in V\backslash\{r\}$, $\mathsf{LB}(q):=(2/Q)\sum_{i\neq r}q_i\cdot d(r,i)$ is a lower bound on the optimal solution length (see, e.g., Haimovich and Rinnooy Kan 1985). Hence, the expected value $E[\mathsf{LB}(q)]=(2/Q)\sum_{i\neq r}E[\xi_i]\cdot d(r,i)$ is a lower bound for VRPSD. These lower bounds were also used in Bertsimas (1992).

3. Split-Delivery VRPSD

In this section we consider the split-delivery problem, where the demand at any location may be satisfied over multiple visits to that location. We first present the randomized approximation algorithm. The vehicle visits vertices in the order given by any α-approximate TSP tour. The main idea in our algorithm is to initially fill the vehicle with a *random number of items* (and later, whenever the vehicle is empty, perform refills to maximum capacity Q). This might seem wasteful compared to the more natural strategy of deterministically filling up to maximum capacity Q; however, as shown in the following example, the expected vehicle route under our algorithm can be much shorter. Indeed, this is precisely the reason why we can obtain a better approximation ratio for our randomized algorithm (Theorem 3.1) than for the deterministic algorithm (Theorem 3.2).

Example. Consider a VRP instance on four locations $\{r,a,b,c\}$, with r being the depot. There is a demand of $Q-1$ at each of a and c, and demand of 2 at b; recall that Q is the vehicle capacity. The TSP tour is $\tau=(r,a,b,c,r)$ with distances $d(r,a)=d(c,r)=1$ and $d(a,b)=d(b,c)=L$ where $L\gg 1$ is a large value. The metric is given by shortest-path distances on τ. See also Figure 1.

If the vehicle starts with being filled deterministically to Q, it needs to refill from b and the route followed is $\sigma_d:=(r,a,b,r,b,c,r)$. This has length $d(\sigma_d)=d(\tau)+2\cdot d(r,b)=4(L+1)$.

On the other hand, suppose the vehicle is initialized with a uniformly random quantity $\ell\in\{0,\ldots,Q\}$. In this case, there are four possible vehicle routes shown below (because demands at some vertices may be zero, the vehicle performs a refill trip only when it is empty *and* at a vertex of positive residual demand).

1. $\ell=Q$. The vehicle route is just σ_d, of length $4(L+1)$.

Figure 1. VRP instance with TSP tour $\tau = (r, a, b, c, r)$.

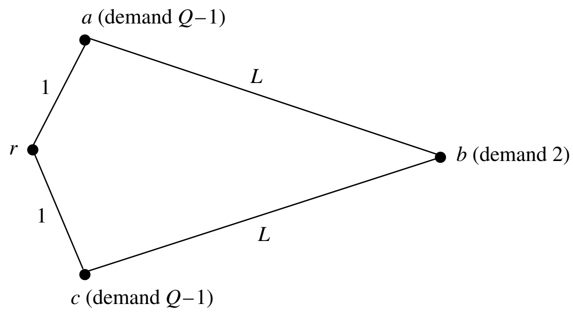

2. $\ell = Q - 1$. The vehicle refills from b and c, so the route is $\sigma_r^1 := (r, a, b, r, b, c, r, c, r)$, of length $4L+6$.
3. $1 \leqslant \ell \leqslant Q-2$. The vehicle refills from a and c, so the route is $\sigma_r^2 := (r, a, r, a, b, c, r, c, r)$, of length $2L+6$.
4. $\ell = 0$. The vehicle refills from a and b, and the route is $\sigma_r^3 := (r, a, r, a, b, r, b, c, r)$, of length $4L+6$.

Now the expected length of the vehicle route equals:

$$\frac{1}{Q+1} \cdot d(\sigma_d) + \frac{1}{Q+1} \cdot d(\sigma_r^1) + \frac{Q-2}{Q+1} \cdot d(\sigma_r^2) + \frac{1}{Q+1} \cdot d(\sigma_r^3) = 2L + 6 + \frac{6L-2}{Q+1}.$$

As $Q \to \infty$, this expectation tends to $2L+6$, which is smaller than the deterministic route length of $4L+4$.

THEOREM 3.1. There is a randomized $(1+\alpha)$-approximation algorithm for VRPSD with split deliveries.

PROOF. The algorithm SplitALG proceeds as follows.

1. Compute an α-approximate TSP tour τ on all vertices.
2. Number the vertices such that r is numbered 0 and tour τ visits vertices in the order $0, 1, 2, \ldots, |V|-1$.
3. Choose a uniformly random value $\ell \in [0, Q]$, and initialize the vehicle to carry ℓ units.
4. The vehicle starts from vertex 0 (the depot), and for each $i = 1, \ldots, |V|-1$ does the following:
   (a) Let $\tilde{Q}_i$ be the units (of the item) carried by the vehicle when it visits vertex i (note that $\tilde{Q}_1 = \ell$).
   (b) Let q_i be the demand observed at vertex i.
   (c) If $q_i \leqslant \tilde{Q}_i$, then serve the demand at i and move on to vertex $i+1$ (with $\tilde{Q}_{i+1} \leftarrow \tilde{Q}_i - q_i$).
   (d) If $q_i > \tilde{Q}_i$, then serve $\tilde{Q}_i$ units of demand at i and make a refill visit to and from r:
   - The vehicle fills up Q units at r, returns to i and serves the remaining $q_i - \tilde{Q}_i$ demand at i.

Then continue on to vertex $i+1$ with $\tilde{Q}_{i+1} \leftarrow Q - (q_i - \tilde{Q}_i)$.

We will bound the expected length of the solution obtained by the above algorithm. In the analysis, we first *condition* on the realization q_i of demands at all vertices $i \in V \setminus \{r\}$. The initial vehicle load ℓ at the depot r remains uniformly random in $[0, Q]$. A vertex $i \in V$ is called a *breakpoint* if the vehicle executes a refill trip to r from i (i.e., Step 4(d) applies at vertex i). This happens precisely when the vehicle becomes empty at vertex i while there is still unserved demand at i. Observe that vertex i is a breakpoint iff there is $p \in \mathbb{Z}_{\geqslant 0}$ such that $\sum_{j=1}^{i-1} q_j \leqslant \ell + p \cdot Q < q_i + \sum_{j=1}^{i-1} q_j$. Because ℓ is the only random variable, and it is uniform in $[0, Q]$, we have:

$$\Pr[i \text{ is break point}] = \frac{q_i}{Q}.$$

The solution length conditioned on demands $\{q_i\}_{i \in V\setminus\{r\}}$ equals $d(\tau) + 2\sum_{i \neq r} d(r, i) \cdot \mathbb{I}(i \text{ breakpoint})$, where $d(\tau)$ is the length of the tour τ under the length function d, and $\mathbb{I}(i \text{ breakpoint})$ is the indicator random variable for the event that i is a break point (for each $i \neq r$). Hence, the expected solution length (conditioned on the demands $\{q_i\}_{i \in V\setminus\{r\}}$) is:

$$d(\tau) + 2\sum_{i \neq r} \Pr[i \text{ is break point}] \cdot d(r, i) = d(\tau) + \frac{2}{Q} \sum_{i \neq r} q_i \cdot d(r, i) = d(\tau) + \mathsf{LB}(q)$$

Recall that this expectation is only over the choice of ℓ. Now, taking expectations over the demands, the expected solution length equals $d(\tau) + E[\mathsf{LB}(q)]$. Because τ is an α-approximate minimum TSP tour and $E[\mathsf{LB}(q)]$ is a lower bound for VRPSD, the expected length of this solution is at most $(1+\alpha)$ times the optimal value of the VRPSD instance. □

We now consider the algorithm SplitALG if we start with deterministic value $\ell = Q$ (vehicle is initially filled to capacity). We show next that even this algorithm achieves a constant approximation guarantee. This algorithm performs no better than the *cyclic heuristic* suggested in Bertsimas (1992), and hence we establish the same upper bound for the cyclic heuristic. It was shown (Bertsimas 1992) that the cyclic heuristic achieves a worst-case guarantee of $\alpha + 1 + o(1)$ in the case of *identical demand distributions* at all vertices, and a $Q + \alpha$ guarantee for general demands (Bertsimas 1992, Theorem 4). Moreover, it was conjectured that the worst-case guarantee of the cyclic heuristic is a constant (independent of Q) even with general demand distributions. The following theorem shows that this is indeed the case.

THEOREM 3.2. The cyclic heuristic is a $(1+2\alpha)$-approximation algorithm for VRPSD.

PROOF. Let τ denote an α-approximate TSP tour. $0, 1, 2, \ldots, |V|-1$ where r is numbered 0. The cyclic heuristic (Bertsimas 1992) considers $|V|-1$ different strategies: obtained by visiting vertices in the order $\langle 0, i, i+1, \ldots, i+|V|-1 \rangle$ modulo $|V|$ for each $1 \leqslant i \leqslant |V|-1$, and returns the strategy with the least objective value. We prove the claimed upper bound for even the (possibly weaker) algorithm that only considers the vertex ordering $0, 1, 2, \ldots, |V|-1$ (as in τ). The vehicle starts from 0, being filled with

Q units, and traverses τ while performing refill trips to the depot whenever it is empty (i.e., algorithm SplitALG with $\ell = Q$). As in the proof of Theorem 3.1, we first condition on the demand realization q_i at each vertex $i \in V \backslash \{r\}$. A vertex $i \in V$ is called a *breakpoint* if the vehicle executes a refill trip to r from i. Let U be the set of all breakpoints and $|U| = u$; for notational uniformity, we also include r as a break point. Then the length of the vehicle's route is $d(\tau) + 2\sum_{w\in U} d(r, w)$. We now establish the following key claim. □

CLAIM 3.3. We have

$$2 \cdot \sum_{w\in U} d(r, w) \leqslant d(\tau) + \frac{2}{Q}\sum_{v\in V} q_v \cdot d(r, v).$$

PROOF. Let the breakpoints U consist of $r = \beta_0, \beta_1, \ldots, \beta_{u-1}$ in that order along τ. For any $l \in \{0, 1, \ldots, u-1\} := [u]$, let τ_l denote the portion of tour τ between vertices β_l and β_{l+1} inclusively (the indices are modulo u). Note that τ is the concatenation $\tau_0 \cdot \tau_1 \cdots \tau_{u-1}$, i.e., $\sum_{l=0}^{u-1} d(\tau_l) = d(\tau)$. For each $l \in [u]$, define a subtour originating and ending at r as $\pi_l := (r, \beta_l) \cdot \tau_l \cdot (\beta_{l+1}, r)$. Observe that the route traced by the vehicle is precisely the concatenation $\pi_0 \cdot \pi_1 \cdot \pi_2 \cdots \pi_{u-1}$. Because the vehicle capacity is Q and it makes refill trips only when it runs out of items, the vehicle delivers exactly Q units in each segment π_l (for $0 \leqslant l \leqslant u-2$), and $Q' \leqslant Q$ units in the last segment π_{u-1}. For each $l \in [u]$ and vertex $i \in \tau_l$, let $C_l(i)$ denote the number of units delivered at vertex i by the vehicle in segment π_l. For technical reasons, set $C_{u-1}(r) := Q - Q'$. From the preceding discussion, we obtain:

$$\sum_{i\in\tau_l} C_l(i) = Q, \quad \forall\, l \in [u], \quad \text{and}$$

$$\sum_{l \mid i\in\tau_l} C_l(i) = q_i, \quad \forall\, i \in V\backslash\{r\}. \tag{1}$$

Consider any fixed segment π_l (for $l \in [u]$). For any vertex $i \in \tau_l$, let $t(i, \beta_l)$ (respectively, $t(i, \beta_{l+1})$) denote the length along τ_l, from β_l to i (respectively, i to β_{l+1}). It follows that $t(i, \beta_l) + t(i, \beta_{l+1}) = d(\tau_l)$ for all vertices $i \in \tau_l$. By the triangle inequality, we have:

$$\left.\begin{aligned} d(\beta_l, r) &\leqslant d(\beta_l, i) + d(i, r) \leqslant t(\beta_l, i) + d(i, r)\\ d(\beta_{l+1}, r) &\leqslant d(\beta_{l+1}, i) + d(i, r) \leqslant t(\beta_{l+1}, i) + d(i, r)\end{aligned}\right\}$$
$$\text{for all vertices } i \in \tau_l.$$

Taking a convex combination of the first (respectively, second) set of inequalities, with multiplier $C_l(i)/Q$ for each $i \in \tau_l$, we obtain the following. (Equation (1) implies that these are indeed convex multipliers.)

$$d(\beta_l, r) \leqslant \sum_{i\in\tau_l} \frac{C_l(i)}{Q} \cdot (t(\beta_l, i) + d(i, r)),$$

$$d(\beta_{l+1}, r) \leqslant \sum_{i\in\tau_l} \frac{C_l(i)}{Q} \cdot (t(\beta_{l+1}, i) + d(i, r)).$$

Adding these two inequalities (using properties of $t(\cdot)$ and $C_l(\cdot)$ from above),

$$\begin{aligned} d(\beta_l, r) + d(\beta_{l+1}, r) &\leqslant \sum_{i\in\tau_l} \frac{C_l(i)}{Q} \cdot (t(\beta_l, i) + t(\beta_{l+1}, i))\\ &\quad + 2\sum_{i\in\tau_l} \frac{C_l(i)}{Q} \cdot d(i, r)\\ &= d(\tau_l)\sum_{i\in\tau_l} \frac{C_l(i)}{Q} + \frac{2}{Q}\sum_{i\in\tau_l} C_l(i) \cdot d(i, r)\\ &= d(\tau_l) + \frac{2}{Q}\sum_{i\in\tau_l} C_l(i) \cdot d(i, r).\end{aligned}$$

Finally, adding the above inequality over $l \in [u]$, where the indices l are modulo u,

$$\begin{aligned} 2 \cdot \sum_{w\in U} d(w, r) &= \sum_{l=0}^{u-1}(d(\beta_l, r) + d(\beta_{l+1}, r))\\ &\leqslant \sum_{l=0}^{u-1} d(\tau_l) + \frac{2}{Q}\sum_{i\in V}\sum_{l \mid i\in\tau_l} C_l(i) \cdot d(i, r)\\ &\leqslant d(\tau) + \frac{2}{Q}\sum_{i\in V\backslash r} q_i \cdot d(i, r).\end{aligned}$$

The last inequality uses Equation (1). Thus we obtain the claim. □

Claim 3.3 gives $2 \cdot \sum_{w\in U} d(r, w) \leqslant d(\tau) + \mathsf{LB}(q)$, which implies that conditioned on demands $\{q_i\}_{i\in V}$, the solution length is at most $2 \cdot d(\tau) + \mathsf{LB}(q)$. Taking expectations, we obtain the desired bound on the cyclic heuristic. □

REMARK. We note that Theorems 3.1 and 3.2 hold even if the range of demand random variables is unbounded. The only difference in the proof of Theorem 3.1 would be to use E[number of refill trips from i] in place of Pr[i is the break point]. For Theorem 3.2, we can treat the breakpoints as a multiset and the same analysis holds.

4. Unsplit-Delivery VRPSD

We now consider the unsplit-delivery problem, where the demand at any location has to be satisfied in a single visit to it, and prove a constant approximation guarantee. This algorithm has one notable difference from the split-delivery algorithm—the vehicle does not necessarily refill to full capacity Q every time it revisits the depot; this aspect seems crucial to proving the performance guarantee.

THEOREM 4.1. There is a randomized $(2+\alpha)$-approximation algorithm for VRPSD with unsplit deliveries.

PROOF. The algorithm UnsplitALG is very similar to the split-delivery case and is given below.

1. Compute an α-approximate TSP tour τ on all locations.

2. Number the locations such that r is numbered 0 and the tour τ visits locations in the order $0, 1, 2, \ldots, |V| - 1$.
3. Choose a uniformly random value $\ell \in [0, Q]$ and initialize the vehicle to carry ℓ units.
4. The vehicle starts at location 0, and for each $i = 1, \ldots, |V| - 1$ does:
   (a) Let $\tilde{U}_i$ be the units (of the item) carried by the vehicle when it visits vertex i.
   (b) Let q_i be the demand observed at vertex i.
   (c) If $q_i \leqslant \tilde{U}_i$, then serve the demand at i and move on to vertex $i+1$ (with $\tilde{U}_{i+1} \leftarrow \tilde{U}_i - q_i$).
   (d) If $q_i > \tilde{U}_i$, then make *two* visits to and from the depot:
   - In the first visit, the vehicle fills up till q_i units at r and serves the demand at i.
   - In the second visit, the vehicle fills up till $Q + \tilde{U}_i - q_i$ units at r, and returns to i.

Then move on to vertex $i+1$ with $\tilde{U}_{i+1} \leftarrow Q + \tilde{U}_i - q_i$.

Just as in Theorem 3.1, we first condition on the realization q_i of demands at all locations $i \in V \backslash \{r\}$. Again, location $i \in V$ is called a *breakpoint* if the vehicle executes a refill trip to r from i (i.e., Step 4(d) applies at location i). We claim that for the same realization of demands, the breakpoints encountered by Algorithms SplitALG and UnsplitALG are *identical*. This follows from the observation that for all locations $i \in V \backslash \{r\}$, the value $\tilde{Q}_i$ (in SplitALG) equals $\tilde{U}_i$ (in UnsplitALG). Thus (from proof of Theorem 3.1), we have $\Pr[i \text{ is breakpoint}] = q_i/Q$ for algorithm UnsplitALG as well. Note that the solution length in UnsplitALG equals $d(\tau) + 4\sum_{i \neq r} d(r, i) \cdot \mathbb{I}(i \text{ breakpoint})$, where $\mathbb{I}(i \text{ breakpoint})$ is the indicator random variable for i being a breakpoint (for each $i \neq r$). Hence, the expected solution length (conditioned on $\{q_i\}_{i \in V \backslash \{r\}}$) is:

$$d(\tau) + 4\sum_{i \neq r} \Pr[i \text{ is the breakpoint}] \cdot d(r, i) = d(\tau) + 2 \cdot \frac{2}{Q} \sum_{i \neq r} q_i \cdot d(r, i) = d(\tau) + 2 \cdot \mathsf{LB}(q)$$

Unconditionally, the expected solution length from UnsplitALG is $d(\tau) + 2 \cdot E[\mathsf{LB}(q)]$. Noting that τ is an α-approximate TSP tour, and $E[\mathsf{LB}(q)]$ is a lower bound for even split-delivery VRPSD, we obtain that UnsplitALG achieves a $(2+\alpha)$ approximation for unsplit-delivery VRPSD.

Remark. A more natural strategy than the randomized one above would be to always refill to capacity Q. We note that this algorithm also achieves a constant approximation guarantee: this can be proved along the lines of Theorem 3.2. However, the resulting bound is again weaker than what is achievable using randomization.

5. Conclusions

In this paper, we presented an $(1+\alpha)$-approximation algorithm for the split-delivery VRP with stochastic demands (where α is the approximation ratio for TSP), and a $(2+\alpha)$-approximation for the unsplit-delivery version of the problem. The natural open question is to improve these guarantees: because the presented bounds match those for the deterministic versions, the first step in this direction would be to improve the approximation ratio for deterministic VRP, which has not seen any real improvement since Haimovich and Rinnooy Kan (1985).

Acknowledgments

A. Gupta was supported in part by NSF award CCF-0729022 and an Alfred P. Sloan Fellowship. R. Ravi was supported in part by NSF grant CCF-0728841.

References

Altinkemer, K., B. Gavish. 1987. Heuristics for unequal weight delivery problems with a fixed error guarantee. *Oper. Res. Lett.* **6**(4) 149–158.
Altinkemer, K., B. Gavish. 1990. Heuristics for delivery problems with constant error guarantees. *Transportation Res.* **24**(4) 294–297.
Arora, S. 1998. Polynomial-time approximation schemes for Euclidean TSP and other geometric problems. *J. ACM* **45**(5) 753–782.
Bertsimas, D. J. 1992. A vehicle routing problem with stochastic demand. *Oper. Res.* **40**(3) 574–585.
Bertsimas, D. J., D. Simchi-Levi. 1996. A new generation of vehicle routing research: Robust algorithms, addressing uncertainty. *Oper. Res.* **44**(2) 286–304.
Bompadre, A., M. Dror, J. B. Orlin. 2006. Improved bounds for vehicle routing solutions. *Discrete Optim.* **3**(4) 299–316.
Bompadre, A., M. Dror, J. B. Orlin. 2007. Probabilistic analysis of unit-demand VRP. *J. Appl. Probab.* **44** 259–278.
Christofides, N. 1977. Worst-case analysis of a new heuristic for the travelling salesman problem. GSIA, CMU-Report 388, Carnegie Mellon University, Pittsburgh.
Das, A., C. Mathieu. 2010. A quasi-polynomial time approximation scheme for Euclidean capacitated vehicle routing. *Proc. Symp. Discrete Algorithms*, ACM-SIAM, Philadelphia.
Dror, M. 2002. Vehicle routing with stochastic demands: Models and computational methods. M. Dror, P. L'Ecuyer, F. Szidarouszky, eds. *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, Vol. 46. Kluwer Academic Publishers, Dordrecht, The Netherlands, 625–649.
Gendreau, M., G. Laporte, R. Séguin. 1996. Stochastic vehicle routing. *Eur. J. Oper. Res.* **88**(1) 3–12.
Haimovich, M., A. H. G. Rinnooy Kan. 1985. Bounds and heuristics for capacitated routing problems. *Math. Oper. Res.* **10**(4) 527–542.
Mitchell, J. S. B. 1999. Guillotine subdivisions approximate polygonal subdivisions: Part II—A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM J. Comput.* **28**(4) 1298–1309.
Powell, W. B., P. Jaillet, A. Odoni. 1995. Stochastic and dynamic networks and routing. M. O. Ball, T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. *Network Routing*. Handbooks in Operations Research and Management Science, Vol. 8. Elsevier Science BV, Amsterdam, 141–295.
Stewart, W., B. Golden. 1983. Stochastic vehicle routing: A comprehensive approach. *Eur. J. Oper. Res.* **14**(4) 371–385.