Appendix A: Mathematical Proofs

In this appendix, we provide the proofs of all the mathematical statements that appear in the paper. In order to facilitate the reading experience, we accompany each proof with its corresponding statement as it originally appears in the main document. All statements and proofs are provided in the same order of appearance.

PROPOSITION 1. *A PDTSPN instance that satisfies Assumption 1 admits an optimal tour that visits each neighborhood exactly once. Moreover, such tour is a shortest Hamiltonian tour where the visits take place at some set of points* $\{p_i\}_{i\in I}$, *where* $p_i \in U_i$ *for all* $i \in I$.

Proof. Given the problem's setting, each neighborhood except U_0 has either $d_i > 0$ or $s_i > 0$. Hence, to satisfy all the pickup and delivery requirements, the vehicle needs to visit each neighborhood at least once. For the first claim, all is left to show is that there exists an optimal tour that visits each neighborhood at most once. Notice that the problem is feasible by the first statement in Assumption 1 and the optimality is always obtainable due to the compactness of all the neighborhoods. So, we can start with an arbitrarily selected optimal tour γ that is represented by the sequence $(p_0, p_1, p_2, \cdots, p_n)$, i.e., a closed tour that starts and ends at p_0 and traverses from p_1 to p_n in between, where every point p_k belongs to some neighborhood U_i that may have been visited more than once. We call the number of visits to a neighborhood U_i in this sequence the multiplicity of U_i. If the multiplicity of every neighborhood is one in γ, we are done. Now, suppose otherwise, then the following procedure will reduce by one the multiplicity of any neighborhood that has been visited repeatedly in γ. By applying this process as many times as necessary, one can produce an optimal tour that is Hamiltonian, thus proving the claim.

The procedure is as follows. If a neighborhood that has been visited multiple times, say neighborhood U_i, is a demand neighborhood, then we use p_k and $p_{k'}$ to denote the points in γ that visit U_i the first two times, and let l_k and $l_{k'}$ be the corresponding delivered commodities during each visit. Then, we split the tour γ into several subpaths as

$$\gamma = (p_0 \rightsquigarrow p_{k-1} \to p_k \to p_{k+1} \rightsquigarrow p_{k'} \rightsquigarrow p_n),$$

where "$\to$" means no other neighborhoods are visited in between and "$\rightsquigarrow$" indicates that potentially multiple neighborhoods are visited in between. We construct the following new tour

$$\gamma' = (p_0 \rightsquigarrow p_{k-1} \to p_{k+1} \rightsquigarrow p_{k'} \rightsquigarrow p_n).$$

Let v, v' be two vehicles following tours γ and γ', respectively. Comparing with v, vehicle v' will only take two different actions in tour γ': it skips the stop p_k and it drops $l_k + l_{k'}$ commodities at point $p_{k'}$. By this design, before the point p_{k-1} and after the point $p_{k'}$, the two vehicles are at the same state in their routes, i.e., at each fixed time, they are at the same location with the same amount of commodities. For the tour between p_{k-1} and $p_{k'}$, it is obvious that the second tour is no longer than the first tour (in fact, they have the same length since γ is assumed to be optimal). We only need to show this new construction is possible between p_{k+1} and $p_{k'}$. Notice at the time of reaching p_{k+1}, vehicle v holds strictly less commodities than vehicle v'. Along with the second statement in Assumption 1, it is clear that v' can always replicate the actions that v takes in the subpath between p_{k+1} and $p_{k'}$. Thus, we have reduced the multiplicity of the neighborhood U_i by one in this case. A similar procedure and argument can be applied to the case that the repeatedly visited neighborhood is a supply neighborhood. So, we proved the first claim.

Therefore, we can focus on all the tours that visit each neighborhood exactly once. That consists of all the closed tours that pass through some point set $\{p_i\}_{i\in I}$ with $p_i \in U_i$ for all $i \in I$. Now, because the shortest closed tour for each such fixed point set $\{p_i\}_{i\in I}$ is a Hamiltonian tour, so is the global optimal tour among all the possible point sets. □

LEMMA 1. *If the point p_k is a turning point in an optimal sequence γ^*, then the set of active neighborhoods $\mathcal{U}_k$ is nonempty.*

Proof. $\mathcal{U}_k = \varnothing$ implies either that no neighborhood contains p_k or, if one does, say U_i, then U_i is also intersected by the complement path $\bar{P}_k$. Clearly, in both cases, the sequence $\gamma' = (p_0, \ldots, p_{k-1}, p_{k+1}, \ldots, p_n)$ also visits U_i. Moreover, p_k being a turning point implies that γ' is strictly shorter than the original optimal sequence, which leads to a contradiction. □

THEOREM 1. *If p_k is a turning point in an optimal sequence γ^*, then there must exist a neighborhood $U_i \in \mathcal{U}_k$ such that $U_i \cap P_k = \{p_k\}$. Moreover, if all neighborhoods in $\mathcal{U}_k$ are also path-connected, then there is some $U_i \in \mathcal{U}_k$ such that $U_i \cap T_k = \{p_k\}$.*

Proof. We prove the first statement by contradiction. Let p_k be a turning point such that all neighborhoods $U_i \in \mathcal{U}_k$ intersect with P_k at more points than just p_k. Then, we will construct a strictly shorter feasible sequence γ' that still visits all the neighborhoods, which contradicts the optimality of γ^*. Let p_{k-1} and p_{k+1} be the previous and next points of p_k in γ^*, and let $[p_{k-1}, p_k)$ and $(p_k, p_{k+1}]$ represent the half-open line segment between the two endpoints. For each $U_i \in \mathcal{U}_k$, we define two sets

$$L_i = \arg \max_{p \in U_i \cap [p_{k-1}, p_k)} \delta(p, p_k),$$
$$R_i = \arg \max_{p \in U_i \cap (p_k, p_{k+1}]} \delta(p, p_k).$$

Thus, L_i (or R_i) contains the points in the intersection of U_i and the segment $[p_{k-1}, p_k)$ (or $(p_k, p_{k+1}]$) that is the farthest to point p_k. Notice that, L_i (or R_i) is either a singleton or an empty set, but $L_i \cup R_i$ must be nonempty by the choice of U_i. Then, we define the following two sets

$$L = \arg \min_{p \in \bigcup_{i \in I_k} L_i} \delta(p, p_k),$$
$$R = \arg \min_{p \in \bigcup_{i \in I_k} R_i} \delta(p, p_k),$$

where I_k is the index set of the neighborhoods in $\mathcal{U}_k$. Thus, L (or R) contains the points that are closest to p_k among all the points in the union $\bigcup_{i \in I_k} L_i$ (or $\bigcup_{i \in I_k} R_i$). By the assumption on all U_i's, $L \cup R \neq \varnothing$. Thus, we are left with three possible cases. In the first case, we have $L = \varnothing$ and $R \neq \varnothing$. This means each $U_i \in \mathcal{U}_k$ intersects some point in $(p_k, p_{k+1}]$. Take any $r \in R$, we construct a feasible sequence $\gamma' = (p_0, \ldots, p_{k-1}, r, p_{k+1}, \ldots, p_n)$, which replaces p_k with r. In the second case, $L \neq \varnothing$ but $R = \varnothing$. Similarly, we construct γ' by replacing p_k with l for any $l \in L$. In the last case, both L and R are nonempty. Then, we define $\gamma' = (p_0, \ldots, p_{k-1}, l, r, p_{k+1}, \ldots, p_n)$ for $l \in L$ and $r \in R$. In all three cases, it is easy to check that the new sequence γ' covers all the neighborhoods in $\mathcal{U}_k$ by the construction and is strictly shorter than γ^* by Lemma 2, which contradicts that γ^* is optimal. Hence, we proved the first claim.

For the second statement, suppose every $U_i \in \mathcal{U}_k$ intersects the triangle T_k at some point $q_i \neq p_k$. Then, we can partition the neighborhoods in $\mathcal{U}_k$ into two groups $\mathcal{U}_k^1$ and $\mathcal{U}_k^2$.The former contains neighborhoods that have a path $q_i \rightsquigarrow p_k$ inside the inner region enclosed by the closed curve γ^*; the latter includes all the rest neighborhoods, each of which must have a path $q_i \rightsquigarrow p_k$ intersects the path P_k, because, otherwise, they must have a path $q_i \rightsquigarrow p_k$ intersects the complement path $\bar{P}_k$, which makes U_i non-active by Definition 3, a contradiction. Notice that $\mathcal{U}_k^1$ is nonempty, otherwise, all neighborhoods in $\mathcal{U}_k$ intersect P_k, which contradicts the first statement we have just shown. For each $U_i \in \mathcal{U}_k^1$, let ω_i be a path connecting q_i to p_k and is inside the inner region enclosed by γ^*, and let C_ϵ be a circle of radius ϵ centered at p_k. Then, for a sufficiently small scalar $\epsilon > 0$, the circle C_ϵ intersects all the paths ω_i's as well as the half open line segments $[p_{k-1}, p_k)$ and $(p_k, p_{k+1}]$. We use l and r to denote the intersection points of C_ϵ with $[p_{k-1}, p_k)$ and $(p_k, p_{k+1}]$, respectively. Then, the new sequence $\gamma' = (p_0, \ldots, p_{k-1}, l, r, p_{k+1}, \ldots, p_n)$ visits all the neighborhoods in $\mathcal{U}_k^1$ since the line segment $[l, r]$ intersects all the paths ω_i's. Moreover, for a sufficiently small $\epsilon > 0$, γ' also intersects all neighborhoods in $\mathcal{U}_k^2$, since the intersection points of each neighborhood $U_i \in \mathcal{U}_k^2$ and the path P_k have a fixed positive distance to p_k. Furthermore, the new curve γ' is strictly shorter than γ^* by the construction and the triangle inequality, which contradicts the optimality of γ^*. □

COROLLARY 1. *Let p_k be any turning point of an optimal sequence γ^* and $\mathcal{U}_k$ be the associated active neighborhoods. If all neighborhoods in $\mathcal{U}_k$ are disjoint, then the unique choice $U_i \in \mathcal{U}_k$ satisfies $U_i \cap T_k = \{p_k\}$ for all k.*

Proof. Because all neighborhoods are disjoint, each turning point can belong to exactly one U_i. This shows that $\mathcal{U}_k$ contains only one neighborhood. By the first statement of Theorem 1, $U_i \cap P_k = \{p_k\}$. Suppose U_i also intersects T_k at some point $q_i \neq p_k$. By Lemma 2, replacing q_i with p_k in γ^* produces a strictly shorter feasible sequence, which is a contradiction. □

THEOREM 2. *Given a set of neighborhoods $\mathcal{U}$, a BPC separator H of $\mathcal{U}$, and a visiting sequence σ, every non-redundant local optimal tour γ_σ^* is inside the half-plane H.*

Proof. Let $\mathbf{p} = (p_i)_{i \in I_0}$ be the local optimal curve γ_σ^*, where $p_i \in U_i$ for each $i \in I_0$ and p_i is either a turning or a passing point. Without loss of generality, we assume that the index set I_0 is sorted by the visiting order σ. We use $\hat{\mathbf{p}}$ to denote the corresponding point set $\{p_i\}_{i \in I_0}$. Clearly, if $\hat{\mathbf{p}} \subseteq H$, then the entire tour γ_σ^* also lies within the half-plane H. Suppose otherwise that $\hat{\mathbf{p}} \nsubseteq H$. Then, there are two possible cases: either all or part of the points in $\hat{\mathbf{p}}$ lie in the complement of H.

In the former case, we claim that the projected tour $\text{proj}_H(\mathbf{p}) = (\text{proj}_H(p_i))_{i \in I_0}$ visits every neighborhood U_i in the order of σ and is strictly shorter than γ_σ^*, which induces a contradiction. By assumption, γ_σ^* does not lie in a straight line that parallels to ∂H, which implies the projection of γ_σ^* onto ∂H is strictly shorter than γ_σ^*. It is left to show that $\text{proj}_H(p_i) \in U_i$ for each $i \in I_0$. If the point p_i is on the boundary of U_i, then

$\text{proj}_H(p_i) \in U_i$ is true due to H is a BPC separator of U_i. Otherwise, p_i is an interior point of U_i. Then, the line extended from the line segment $[\text{proj}_H(p_i), p_i]$ intersects U_i at some point $p_i' \in \partial U_i \setminus H$. Then, we have $\text{proj}_H(p_i) = \text{proj}_H(p_i') \in U_i$, where the equality is by the construction of p_i' and the membership relation is due to H is a BPC separator of U_i.

In the latter case, some point $p_0 \in \hat{\mathbf{p}}$ is in the half-plane H. Then, we represent γ_σ^* as a continuous function from the closed interval $[0,1]$ to $\mathbb{R}^2$ such that $\gamma_\sigma^*(0) = \gamma_\sigma^*(1) = p_0$, and $\gamma_\sigma^*(t_i) = p_i$ for all $i \in I_0$. By assumption, there is also some point in $\hat{\mathbf{p}}$ that is not in H, i.e., $\gamma_\sigma^*(t_i) \notin H$ for some $i \in I$, where I is the index set $I_0 \setminus \{0\}$. The fact that γ_σ^* is a continuous function means there exists some closed neighborhood $T = [t_{\min}, t_{\max}] \subseteq [0,1]$ of t_i such that the path $\gamma_\sigma^*(T)$ intersects H only at the two endpoints $p_{\min} = \gamma_\sigma^*(t_{\min})$ and $p_{\max} = \gamma_\sigma^*(t_{\max})$. Let $T' = (t_{\min}, t_{\max})$ be the corresponding open interval of T, and we use the set $I' = \{i \in I \mid t_i \in T'\}$ to label all the neighborhoods that are visited by the path $\gamma_\sigma^*(T')$. Then, we consider the following projected sequence

$$\text{proj}_H(\gamma_\sigma^*(T)) = (p_{\min}, (\text{proj}_H(p_i))_{i \in I'}, p_{\max}).$$

This sequence represents a subtour that visits all the neighborhoods $\{U_i\}_{i \in I'}$ with the same visiting order since H is a BPC separator of $\mathcal{U}$. Also, $\text{proj}_H(\gamma_\sigma^*(T))$ is strictly shorter than the original segment $\gamma_\sigma^*(T)$ since the tour is non-redundant. Moreover, $\text{proj}_H(\gamma_\sigma^*(T))$ and $\gamma_\sigma^*(T)$ shares the same endpoints. Thus, replacing the subtour $\gamma_\sigma^*(T)$ in γ_σ^* with the projected sequence produces a strictly shorter feasible solution, which yields a contradiction. □

Corollary 2. *Given a set of compact neighborhoods $\mathcal{U}$, if the optimal tour γ^* does not lie in a straight line, then the domination hull $D_{\mathcal{U}}$ is a nonempty, convex, compact set.*

Proof. By Theorem 2, given that some optimal curve γ^* is non-redundant, then γ^* lies in $D_{\mathcal{U}}$, which shows $D_{\mathcal{U}}$ is nonempty. $D_{\mathcal{U}}$ is convex and closed since it is the intersection of a set of half-planes. Finally, by definition, a sufficiently large polygon that contains all the neighborhoods can be produced by the intersection of a set of BPC separators. Thus, it also contains $D_{\mathcal{U}}$, which means $D_{\mathcal{U}}$ is bounded, which entails compactness. □

Proposition 2. *Given a set of convex neighborhoods $\mathcal{U}$, the binary search subroutine described in Algorithm 2 converges to the minimized distance*

$$\begin{aligned} \min\ & d \\ s.t.\ & H_{\theta,d} \in \mathcal{H}_{\mathcal{U}}, \\ & d \in [0, r], \end{aligned}$$

when $\epsilon \to 0$.

Proof. Given the initial input $(0, r, \theta, \mathcal{U})$, it suffices to show that, for any value $c \in (0, r)$, if $H_{\theta,c}$ is a BPC separator, then so is $H_{\theta,b}$ for all $b \in (c, r]$; if $H_{\theta,c}$ is not a BPC separator, neither is $H_{\theta,a}$ for all $a \in [0, c)$. To prove the former, we suppose $H_{\theta,c}$ is a BPC separator of $\mathcal{U}$, i.e., it is a BPC separator for every $U \in \mathcal{U}$.

Take any $U \in \mathcal{U}$ and $b \in (c, r]$, the boundary points that are not contained in the half-plane $H_{\theta,b}$ is denoted by $\partial U \setminus H_{\theta,b}$. Clearly, we have the following

$$\partial U \setminus H_{\theta,b} \subseteq \partial U \setminus H_{\theta,c}.$$

Hence, for every $p \in \partial U \setminus H_{\theta,b}$, the projection $\text{proj}_{H_{\theta,c}}(p)$ is inside U, since $H_{\theta,c}$ is BPC separator of U by assumption. Moreover, because $H_{\theta,b}$ is parallel to $H_{\theta,c}$, the line segment $[p, \text{proj}_{H_{\theta,c}}(p)]$ intersects the line $\partial H_{\theta,b}$ at the projection point $\text{proj}_{H_{\theta,b}}(p)$. Thus, $\text{proj}_{H_{\theta,b}}(p) \in U$ due to the convexity of U, which shows $H_{\theta,b}$ is a BPC separator of U by definition. For the latter statement, we take any $a \in [0, c)$ and assume $H_{\theta,c}$ is not a BPC separator of $\mathcal{U}$. That is, for some $U \in \mathcal{U}$, $H_{\theta,c}$ is not a BPC separator of U. This means for some point $p \in \partial U \setminus H_{\theta,c}$, the projection $\text{proj}_{H_{\theta,c}}(p)$ is not in U. Same as before, the point p belongs to $\partial U \setminus H_{\theta,a}$. Towards a contradiction, suppose $H_{\theta,a}$ is a BPC separator of $\mathcal{U}$, then $\text{proj}_{H_{\theta,a}}(p) \in U$. Same as before, the segment $[p, \text{proj}_{H_{\theta,a}}(p)]$ intersects the line $\partial H_{\theta,c}$ at the projection point $\text{proj}_{H_{\theta,c}}(p)$. Then, the convexity of U implies $\text{proj}_{H_{\theta,c}}(p) \in U$, which contradicts the choice of p. This concludes the proof. □

PROPOSITION 3. *Given non-convex region representations in* (3f)*, the cut generation implementation of* (5) *with GBCs from* (7) *solves for an optimal solution of* (3)*.*

Proof. Every $u \in E_{sd}$ represents a feasible sequence of neighborhoods. Hence, after solving the corresponding subproblem $\Phi(u)$ optimally, the right-hand-side of new constraint (7) enforces the actual length of the shortest Hamiltonian tour under the given sequence u. Therefore, regardless of the convexity of (3f), the GBCs from (7) yield a sharp underestimator θ of the tour lenght, which guarantees the correctness of the cut generation method from the classic upper and lower bounding argument of the Benders decomposition. □