LoRaPCR: Long Range Point Cloud Registration through Multi-hop Relays in VANETs

Zhenxi Wang[1], Hongzi Zhu[1*], Yunxiang Cai[1], Quan Liu[1], Shan Chang[2], Liang Zhang[1]

[1]*Shanghai Jiao Tong University, Shanghai, China*

[2]*Donghua University, Shanghai, China*

{leslie1956, hongzi, caiyunxiang, liuquan2017, zhangliang}@sjtu.edu.cn, changshan@dhu.edu.cn

Abstract—**Point cloud registration (PCR) can significantly extend the visual field and enhance the point density on distant objects, thereby improving driving safety. However, it is very challenging for vehicles to perform online registration between long-range point clouds. In this paper, we propose an online long-range PCR scheme in VANETs, called LoRaPCR, where vehicles achieve long-range registration through multi-hop short-range highly-accurate registrations. Given the NP-hardness of the problem, a heuristic algorithm is developed to determine best registration paths while leveraging the reuse of registration results to reduce computation costs. Moreover, we utilize an optimized dynamic programming algorithm to determine the transmission routes while minimizing the communication overhead. Results of extensive simulations demonstrate that LoRaPCR can achieve high PCR accuracy with low relative translation and rotation errors of 0.55 meters and 1.43°, respectively, at a distance of over 100 meters, and reduce the computation overhead by more than 50% compared to the state-of-the-art method.**

Index Terms—**point cloud registration, cooperative sensing, multi-hop relay, VANETs**

I. INTRODUCTION

Due to its precise panoramic view and nighttime adaptability, Light Detection and Ranging (LiDAR) sensors are installed on newly manufactured vehicles to perform tasks such as object detection [1] [2] and semantic segmentation [3], significantly improving driving safety. However, LiDAR is also accompanied by insufficient long-range resolution as illustrated in Figure 1(a) and occlusion by obstacles as illustrated in Figure 1(b). With vehicle-to-vehicle (V2V) communications in the Vehicular ad hoc networks (VANETs), sharing and registering point clouds between vehicles become possible [4], which enhances the perception capability of vehicles [5] [6]. For example, Figure 1(c) demonstrates that a car, which is hard to recognize in each individual point cloud, can be identified after two of such point clouds are registered. We refer to the *online long-range point cloud registration (PCR) problem* as the problem that distant vehicles can accurately register their point clouds through one- or multi-hop V2V communication. Promisingly, if long-range point clouds of significant difference in terms of point cloud density and perspective can be well registered, it will greatly improve the performance of downstream tasks.

A feasible scheme for online long-range PCR problem in a vehicular scenario, however, has to meet four rigid

(a) Vehicles in point cloud

(b) Blind zone due to obstacles

(c) Long-range PCR enhances vehicular perception capability
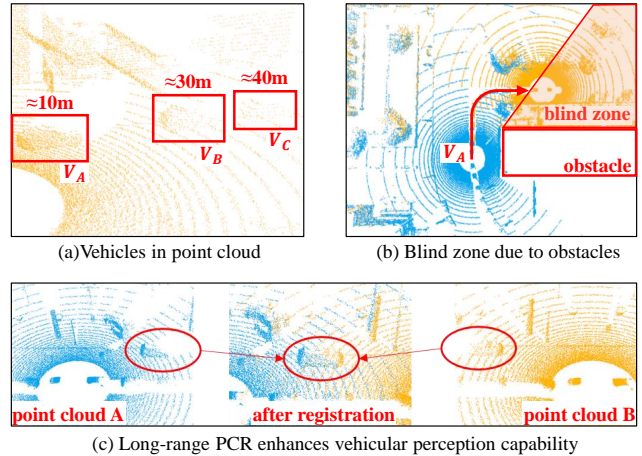
Fig. 1: Motivation of long-range PCR. (a) The resolution of a point cloud decreases with the distance to the LiDAR sensor. (b) The perception field of LiDAR is blocked by obstacles. (c) Point clouds collected on two vehicles with different viewpoints can be registered (aligned), enhancing the perception capability of both vehicles.

requirements as follows: 1) the scheme should be able to deal with long-range registration (*e.g.*, tens of meters) to facilitate downstream perception tasks; 2) due to the rapid movement of vehicles, the scheme should achieve online registration with low response time to ensure the timeliness of the registration results; 3) given the limited resources on a vehicle, such a scheme should be efficient in terms of computation and communication costs; 4) such a scheme should achieve high accuracy as the registered point clouds may be utilized in critical driving safety applications.

In the literature, most PCR work focuses on short-range point cloud alignment (*e.g.*, within 10 meters), where they register point clouds with similar density and significant overlap. In these methods, feature extractors aim to improve the quality of extracted representations, while outlier rejection methods [7] [8] attempt to distinguish erroneous responses. However, these methods fail to address long-range PCR challenges. While [6] [9] address the long-range problem by training a distance-insensitive feature extractor, these approaches lack optimization for communication overhead. EMP [5] enhances vehicles using point clouds from roadside infrastructure, but it does not consider utilizing multi-hop V2V communication

to achieve long-range registration. As a result, to the best of our knowledge, there is no cost-efficient scheme which can accurately register long-distance point clouds in vehicular settings.

In this paper, we propose *LoRaPCR*, an online long-range PCR scheme particularly designed for VANETs. The main idea of LoRaPCR is to utilize previous short-range PCRs of high accuracy to establish a superior multi-hop registration path for present long-range PCRs between each pair of vehicles. To this end, a base station (BS) or a roadside unit (RSU) first collects the information about the location, previous PCR records, and the current PCR requests of vehicles in its vicinity, and then makes the best multi-hop registration and transmission strategies, with the goal of guaranteeing a low accumulative PCR errors while optimizing the global computation and communication overheads.

Two main challenges are encountered. First, determining registration paths for the registration requests is challenging. Due to its NP-hardness and the enormous search space, obtaining an optimal solution in an online system is difficult. To address this challenge, we propose a heuristic algorithm based on the shortest path breadth-first search (BFS). More specifically, LoRaPCR prioritizes the simplest and most efficient one-hop requests. Then, when dealing with hard requests, it maximizes the reuse of previously used PCRs to reduce computational costs. For each hard request, we utilize a queue-optimized BFS algorithm to accelerate the search process and ensure the discovery of all feasible solutions, achieving the highest request satisfaction ratio.

Second, it is challenging to determine globally optimal transmission routes for registration requests. This issue is equivalent to the Graphical Steiner Minimal Tree (GSMT) problem [10], also known to be NP-hard. Although the scale of this problem is smaller than the previous one, existing solutions are far from meeting online requirements. To tackle this challenge, we propose a novel reduction algorithm based on dynamic programming. In this algorithm, we leverage the locality of the communication network to reduce the problem's scale. Specifically, we utilize the connectivity between target nodes to simplify the problem's scale, reducing it from the number of target nodes to the number of connected components. Moreover, we apply pruning techniques for two common special cases to further enhance the algorithm efficiency.

We generate intensive vehicle traces in three typical driving scenarios (*i.e.*, town, highway, and village) using the CARLA simulator [11] to validate the performance of LoRaPCR. We examine three different registration request preferences (*i.e.*, near, far, and uniform) and observe an average request satisfaction ratio (RSR) of approximately 66% under the most challenging preference (*i.e.*, far). In comparison, the state-of-the-art (SOTA) method achieved an RSR of only 4%. Extensive simulations are conducted and the results demonstrate that LoRaPCR significantly reduces the computation overhead by more than 50% compared to the SOTA method. To the best of our knowledge, LoRaPCR is the first PCR system in VANETs that achieves registration with point clouds separated
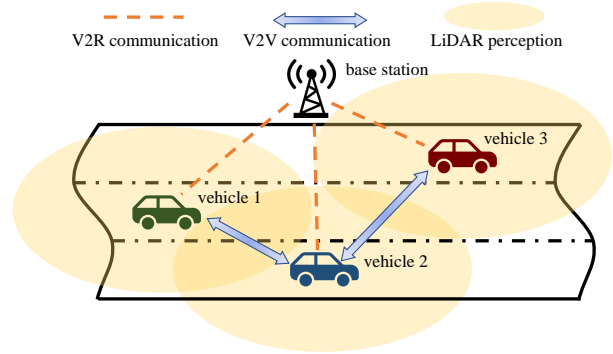

Fig. 2: Illustration of the system model

over 100 meters, obtaining a high PCR accuracy with low relative translation and rotation errors of 0.55 meters and 1.43°, respectively.

We highlight the main contributions made in this work as follows: 1) A multi-hop PCR path selection algorithm based on a heuristic shortest path approach is proposed, effectively reducing the computational costs; 2) An online optimal transmission path selection algorithm is proposed, integrating locality-based reduction and the dynamic programming techniques; 3) extensive experiments are conducted and results demonstrate the efficacy of LoRaPCR, achieving SOTA performance for long-range PCRs.

II. System Model and Problem Definition

A. System Model

In the long-range PCR problem, we consider the following two types of entities as shown in Figure2.

- **Vehicles:** are peers with equal capabilities as follows: 1) vehicles are capable of performing neural network computations to accomplish PCR tasks; 2) vehicles are capable of high-speed (*e.g.*, more than 750Mb/s) V2V communication with a communication range of tens of meters, such as via millimeter-wave (mmWave). Additionally, vehicles can communicate with a BS or an RSU via longe-range communication links; 3) vehicles are equipped with LiDAR sensors that are capable of generating and storing point cloud data; 4) vehicles can obtain rough distance estimations of nearby vehicles but do not require precise localization. During online registration, we do not require any side channel information of vehicle position.
- **BS:** has the capability to utilize vehicle-to-roadside (V2R) communication for obtaining the information about the location of nearby vehicles, previous registration results and new PCR requests. The BS has basic computation and storage capabilities to make PCR strategies, which can be sent to vehicles for execution.

B. Problem Definition

We denote the sequence $\{X_1^{v_i}, X_2^{v_i}, ..., X_t^{v_i}\}$ as the time series of t point cloud frames of vehicle $v_i \in V$, where
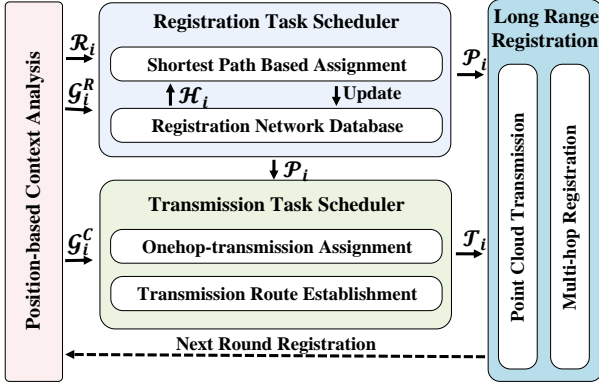
Fig. 3: The system architecture of LoRaPCR

$X_k^{v_i} = \{p_n \in \mathbb{R}^3 | n = 1, 2, ..., N\}$ for $k \in [1, t]$ is the k-th point cloud frame of N points. At time t, vehicle v_i generates registration requests $\mathcal{R}_t^{v_i} \subseteq \{(v_i, v_j) | v_j \in I^{v_i}\}$, where I^{v_i} represents the interest region of v_i. A PCR request (v_i, v_j) is satisfied if there is a registration result (*i.e.*, the transform matrix $M_{i,j}$ that aligns $X_t^{v_j} M_{i,j}^T$ with $X_t^{v_i}$) and the target point cloud data $X_t^{v_j}$ is transmitted to v_i. For registration request set $\mathcal{R}_t = \bigcup_{v_i \in V} \mathcal{R}_t^{v_i}$, the long-range multi-vehicle PCR problem is to find registration paths $\mathcal{P}_t$ and transmission paths $\mathcal{T}_t$ such that after they are executed, the registration requests are satisfied while the resulting computational and communication costs are minimized. The computational cost is determined by the number of registration tasks required, while the communication cost is determined by the number of transmissions needed along the transmission routes to satisfy each registration request.

III. Design of LoRaPCR

A. Overview

The core idea of LoRaPCR is to achieve long-range PCR through short-range PCR. Instead of rejecting long-range registration requests or obtaining an unreliable registration result, LoRaPCR searches for a high-quality registration path $X_t^{v_i} \leftrightarrow X_{t'}^{v_k} \leftrightarrow ... \leftrightarrow X_t^{v_j}$ to satisfy the long-range registration requests. The relay nodes in the registration path can be nearby point clouds of other vehicles $(X_t^{v_k})$ or even nearby point clouds from previous time instances $(X_{t'}^{v_k})$.

In each round of registration, each vehicle reports the neighbor information to the base station, enabling the base station to construct the registration graph $\mathcal{G}_t^R$ and communication graph $\mathcal{G}_t^C$. Then base station determines the registration paths $\mathcal{P}_t$ and transmission routes $\mathcal{T}_t$ and decomposes them into individual registration tasks and transmission tasks for each vehicle. After every vehicle finishes its respective registration tasks and transmission tasks, the registration requests are satisfied, which means $X_t^{v_j}$ is transmitted to v_i, and v_i has $M_{i,j}$ to align $X_t^{v_j} M_{i,j}$ with $X_t^{v_i}$ for every $(v_i, v_j) \in \mathcal{R}_t$. As depicted in Figure 3, LoRaPCR consists of two main components as follows:


Fig. 4: Illustration of the integrated registration graph $\mathcal{I}_t$, which is integrated by the registration graph $\mathcal{G}_t^R$ and the registration history $\mathcal{H}_t$

Registration Task Scheduler. The registration task scheduler takes the registration requests and the registration graph as input, then combine the registration graph with the registration history and utilizes a shortest-path-based search algorithm to find registration paths for the registration requests. The registration task scheduler then sends the registration paths to the transmission task scheduler as well as vehicles.

Transmission Task Scheduler. The transmission task scheduler decomposes the registration paths into transmission tasks and searches for transmission routes for registration requests within the communication graph. Then the transmission task scheduler sends the transmission routes to vehicles.

B. Registration Task Scheduler

The registration task scheduler is used to determine registration paths $\mathcal{P}_t = \{Path(X_t^{v_i}, X_t^{v_j}) | (v_i, v_j) \in \mathcal{R}_t\}$ and it achieves this goal by three steps:

1) Registration Graph Integration: As shown in Figure4, in the registration graph $\mathcal{G}_t^R$ and registration history $\mathcal{H}_t$, each node denotes a point cloud $X_t^{v_i}$ while each edge $(X_t^{v_i}, X_{t'}^{v_j})$ denotes a potential registration (*e.g.*, in $\mathcal{G}_t^R$) or a completed registration (*e.g.*, in $\mathcal{H}_t$). We define the weight of edges as the computation cost, *i.e.*, $w_{edge} = 1$ for a potential registration while $w_{edge} = 0$ for a completed registration. We can estimate the quality of registration using an empirical model and define it as the attribute of the corresponding edge. The registration task scheduler combines $\mathcal{G}_t^R$ and $\mathcal{H}_t$ into one integrated registration graph $\mathcal{I}_t$, and connect $\mathcal{G}_t^R$ and $\mathcal{H}_t$ by the temporal edges like $(X_t^{v_i}, X_{t'}^{v_i})$. The temporal edge $(X_t^{v_i}, X_{t'}^{v_i})$ means v_i can perform a registration computation between those two point clouds generated and stored by itself. The integrated registration graph represents all available registration edges to find registration paths for registration requests.

2) Registration Path Solution: After obtaining the integrated registration graph, the registration task scheduler determines the Minimum Request Satisfaction Graph(MRSG) for $\mathcal{R}_t$ on $\mathcal{I}_t$. The MRSG is defined as a subgraph with minimal computational cost, in which each request can identify a path if there exists a path in the original graph. The computational cost of MRSG equals the computational cost to implement $\mathcal{P}_t$, which can be represented as $\sum w_{edge}(edge \in \bigcup Path, Path \in \mathcal{P}_t)$. By solving the MRSG problem, the registration task scheduler can derive $\mathcal{P}_t$, which is the output of this module.

Algorithm 1: BFS-based Algorithm for MRSG Problem

Input : $\mathcal{R}_t$, $\mathcal{I}_t$
Output: $\mathcal{P}_t$

1 Initialize $\mathcal{P}_t$ as an empty set
2 **for** *each* $X_t^{v_i}$ *in* $\mathcal{I}_t$ **do**
3     $path\_lists[X_t^{v_i}]$ = ConstrainedSPFA($X_t^{v_i}$, $\mathcal{I}_t$)
4   **for** *each* $X_t^{v_j}$ *such that* $(X_t^{v_i}, X_t^{v_j}) \in \mathcal{R}_t$ **do**
5         $Path(X_t^{v_i}, X_t^{v_j})$ is minimum cost path in $path\_lists[X_t^{v_i}][X_t^{v_j}]$
6   Add $Path(X_t^{v_i}, X_t^{v_j})$ to $\mathcal{P}_t$

7 **return** $\mathcal{P}_t$
8 **Function** ConstrainedSPFA(X_{start}, $\mathcal{I}_t$)
9   Initialize a list $path\_lists[X]$ for every node X in $\mathcal{I}_t$
10   Initialize a queue and add X_{start} to it
11   **while** *the queue is not empty* **do**
12   Take the front node X_f out of the queue
13   **for** *each neighbor* X_n *of* X_f **do**
14   Try to add new path to $path\_lists[X_n]$
15   **if** *quality of new path* $\geq$ *quality of any path in* $path\_lists[X_n]$ *and cost of new path* $<$ *cost of that path* **then**
16   Update the $path\_lists[X_n]$
17   If X_n is not in the queue, add it to the queue

18   **return** $path\_lists$

3) Registration Database Update: The registration task scheduler updates $\mathcal{H}_t \rightarrow \mathcal{H}_{t+1}$ according to $\mathcal{P}_t$ and the registration results returned by vehicles.

The MRSG problem is NP-hard, which can be proven by establishing a reduction from the Graphical Steiner Minimal Tree (GSMT) problem [10] with unit weight, as known to be NP-hard. To address the MRSG problem, one of the most straightforward approximate algorithms is to use the BFS method to find a shortest path for each request. BFS can achieve a high request satisfaction ratio by considering every possible path for each request, but it may not achieve the global minimum cost as it only considers local optimality. The pseudocode for the BFS-based algorithm is shown in Algorithm 1.

Considering the practical application requirements, we add the following two constraints to the registration paths:

- The registration recall of the registration path, which is the product of the registration recalls of each edge, must exceed a lower bound of 95%.
- The number of hops in the registration path is limited by an upper bound, which restricts the error range of the registration result obtained from the registration path.

Specifically, we employ the Shortest Path Faster Algorithm (SPFA), a queue-optimized BFS algorithm, to solve MRSG problem, given that the constraint conditions prevent infinite loops.

One way to optimize the BFS algorithm is by utilizing the overlap between different paths to reuse registration edges as much as possible. A feasible approach is to record the edges used in the determined paths and to reuse them as much as possible. In this approach, the order in which the paths for

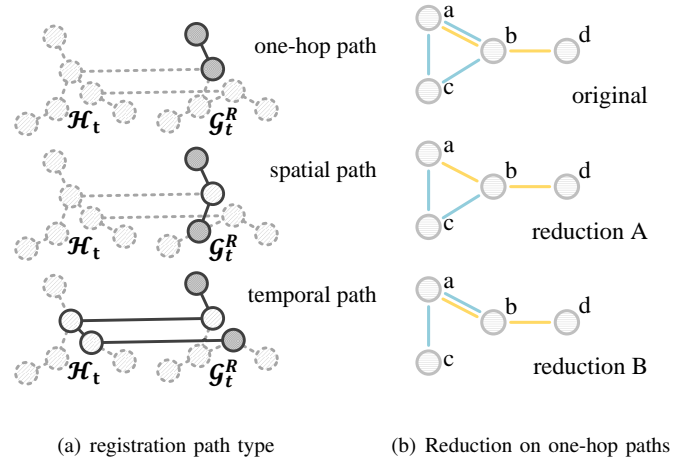(a) registration path type   (b) Reduction on one-hop paths

Fig. 5: Illustration of registration paths.(a)Three types of registration paths.(b)Effective reductions on one-hop paths can decrease the number of registrations.

different requests are determined becomes crucial to achieve efficient path reuse. As shown in Figure5(a), the registration paths can be categorized into three types: 1) the one-hop path, which indicates a single hop registration 2) the spatial path, which involves using extra nodes at the same time instance as a relay. 3) the temporal path, which involves using extra nodes from a previous time instance as a relay. Among the three types of paths, one-hop paths are the most efficient and intuitive. If a request has an one-hop path available, then there is high probability that the registration path for this request in the global optimal solution will indeed be an one-hop path. Hence, we propose a novel algorithm called One-Hop Request First (OHRF). The core idea of the OHRF algorithm is to reuse the one-hop paths when making decisions for other requests. More specifically, we first check whether each registration request can be satisfied by a one-hop path. If so, we directly use the one-hop path as the solution. Then, we create a duplicate graph, denoted as $\mathcal{I}'_t$, and set the weight of all edges used in the one-hop solutions to 0 in $\mathcal{I}'_t$. For the other requests, referred to as hard requests, we employ the BFS in $\mathcal{I}'_t$ to search for the shortest path as their solutions. By adopting this approach, we effectively reuse one hop paths when searching for the shortest path for hard requests, thereby saving computational costs.

In addition, we discover that there is also a substantial amount of redundancy among one-hop paths. As depicted in Figure5(b), the blue edges represent paths corresponding to one-hop requests, while the orange edges represent paths corresponding to hard requests. In original case, the hard request reuses one edge from the one-hop request, resulting in a global cost of 4. Under the premise of ensuring that requests are satisfied, we can perform reduction on one hop paths. Reduction A removes some edges while still ensuring the satisfaction of all requests. Unfortunately, this approach may not necessarily reduce the global computation cost, as

Algorithm 2: OHRF-based Algorithm for MRSG Problem with Reduction

---
Input : $\mathcal{R}_t, \mathcal{I}_t$
Output: $\mathcal{P}_t$

1 Initialize $\mathcal{P}_t$ as an empty set
2 Initialize $hard\_requests$ as an empty list
3 Initialize one_hop_paths as an empty list
4 **for** each $(X_t^{v_i}, X_t^{v_j})$ in $\mathcal{R}_t$ **do**
5   **if** $(X_t^{v_i}, X_t^{v_j})$ is one-hop reachable in $\mathcal{I}_t$ **then**
6        $Path(X_t^{v_i}, X_t^{v_j}) = (X_t^{v_i}, X_t^{v_j})$
7   Add $Path(X_t^{v_i}, X_t^{v_j})$ to one_hop_paths
8   Set cost of edge $(X_t^{v_i}, X_t^{v_j})$ in $\mathcal{I}_t$ to 0
9   **else**
10   Add $(X_t^{v_i}, X_t^{v_j})$ to $hard\_requests$

11 **for** each $X_t^{v_i}$ in $\mathcal{I}_t$ **do**
12    $path_lists[X_t^{v_i}] = $ ConstrainedSPFA$(X_t^{v_i}, \mathcal{I}_t)$
13   **for** each $(X_t^{v_i}, X_t^{v_j})$ in $hard\_requests$ **do**
14       $Path(X_t^{v_i}, X_t^{v_j})$ is minimum cost path in $path\_lists[X_t^{v_i}][X_t^{v_j}]$
15   Mark the reused path in one_hop_paths
16   Add $Path(X_t^{v_i}, X_t^{v_j})$ to $\mathcal{P}_t$

17 Construct a graph, G_{hop}, using one_hop_paths
18 Compute the MST, T_{hop}, of G_{hop} while trying to retain marked paths
19 **for** $Path(X_t^{v_i}, X_t^{v_j})$ in one_hop_paths **do**
20   **if** $Path(X_t^{v_i}, X_t^{v_j})$ is in the MST T_{hop} **then**
21   Add $Path(X_t^{v_i}, X_t^{v_j})$ to $\mathcal{P}_t$
22   **else**
23   Set cost of edge $(X_t^{v_i}, X_t^{v_j})$ in $\mathcal{I}_t$ to 0
24       $path_lists[X_t^{v_i}] = $ ConstrainedSPFA$(X_t^{v_i}, \mathcal{I}_t)$
25       $Path(X_t^{v_i}, X_t^{v_j})$ is minimum cost path in $path\_lists[X_t^{v_i}][X_t^{v_j}]$
26   Add $Path(X_t^{v_i}, X_t^{v_j})$ to $\mathcal{P}_t$

27 **return** $\mathcal{P}_t$

Fig. 6: The procedure of decomposing registration path into transmission tasks

(a) MTG   (b) GSMT   (c) Connected components   (d) Reduced GSMT

Fig. 7: The procedure of reducing MTG problem

the removed edges could be reintroduced in the paths of hard requests.

Based on the above observations, we propose a novel reduction method. The core idea is to record edges used by hard requests and constructs the minimum spanning tree to reduce one-hop paths while preserving the recorded edges. It can be seen that in Reduction B the global computation cost is 3. The pseudocode for the OHRF-based algorithm with reduce is shown in Algorithm 2.

C. Transmission Task Scheduler

The transmission task scheduler is used to determine transmission paths $\mathcal{T}_t$ and it achieves this goal by two steps:

1) Transmission Task Decomposition: The transmission task scheduler takes $\mathcal{P}_t$ as input, and for every $Path \in \mathcal{P}_t$, it can be decomposed into several transmission tasks. As shown in Figure 6, $Path(X_t^{v_C}, X_t^{v_A})$ consists of 4 edges. For edge ①, there is a transmission task of v_A transmitting $X_t^{v_A}$ to v_B. For edge ② and ④, there is no transmission task since the point cloud pair belongs to the same vehicle. For edge ③, there is no transmission task as the registration has already been completed at time t'. The completed registration result can be directly obtained from the Registration Network Database.

Besides, there is another transmission task of transmitting $X_t^{v_A}$ to v_C to satisfy the registration request $(X_t^{v_C}, X_t^{v_A})$. Then we can partition the transmission tasks based on the content of the data being transmitted. Each partition consists of a list of transmission tasks of transmitting $X_t^{v_{src}}$ from the source vehicle v_{src} to the target vehicle v_{tgt}.

2) Transmission Tree Solution: After partitioning the transmission tasks, the transmission task scheduler processes each partition sequentially, as they do not interfere with each other. For each partition, the transmission task scheduler needs to determine an optimal transmission route for each transmission task, aiming to minimize the overall communication overhead. This problem is referred to as the Minimum Transmission Graph(MTG) problem and its NP-hardness can be proven by establishing a reduction from the GSMT problem.

For the MTG problem, there exists a dynamic programming algorithm [12] with a time complexity of $O(n^c * 2^m)$, where n represents the number of nodes, and m denotes the number of special nodes in the graph. While we may not be able to reduce the computational complexity, we can improve the algorithm's efficiency by reducing the problem size. For our specific application scenario, we have the following observation: Since the target nodes are all within the interest region of the source node, the distances between the target nodes are not likely to be too long. Consequently, there is a high probability that the target nodes are directly connected in the communication graph. In order to leverage this locality of the communication graph, we try to simplify the GSMT problem by organizing the special nodes into several connected components, which we refer to as "Connected Components Reduction".

As shown in Figure 7(a), there are 5 target vehicles for v_{src}: 2 one-hop transmission task to v_{t_1}, v_{t_2} and 3 transmis-

Algorithm 3: Optimized Dynamic Programming Algorithm with Reduction for GSMT Problem

---

Input : Undirected graph $G = (V, E)$, special node set $Q \subseteq V$
Output: GSMT T
1 Construct subgraph $QG = (Q, QE)$, where
   $QE = \{(q_1, q_2) \mid q_1, q_2 \in Q\}$
2 Find connected components of QG, store MSTs of each connected component
3 Merge nodes within each connected component of QG to form new graph G', Q'
4 **if** *number of connected components in QG = 1* **then**
5   **return** the MST of the single connected component
6 **else if** *number of connected components in QG = 2* **then**
7     $T_{\text{MST1}} \leftarrow$ MST of the first connected component in QG
8     $T_{\text{MST2}} \leftarrow$ MST of the second connected component in QG
9     $u, v \leftarrow$ special nodes representing the two connected components
10     $T_{\text{SP}} \leftarrow$ ShortestPath(G', u, v)
11     $T \leftarrow$ Merge T_{MST1}, T_{MST2}, and T_{SP}
12 **else**
13     $T_{\text{DP}} \leftarrow$ DP(Q', G')
14     $T_{\text{MST}} \leftarrow$ Merge MSTs of all connected components in QG
15     $T \leftarrow$ Merge T_{DP} and T_{MST}
16 **return** T

TABLE I: Ablation Study

Preference	Algorithm	RSR	RCPR	CCPR	RTE(m)	RRE(°)
Far	LoRaBFS	0.66	0.72	1.99	0.53	1.41
	LoRaOHRF	0.66	0.71	**1.99**	**0.52**	**1.40**
	LoRaPCR	**0.66**	**0.49**	2.01	0.53	1.42
Uniform	LoRaBFS	0.81	0.68	1.36	0.38	1.29
	LoRaOHRF	**0.81**	0.56	**1.36**	**0.36**	**1.28**
	LoRaPCR	0.80	**0.35**	1.39	0.39	1.32
Near	LoRaBFS	0.91	0.81	**1.27**	0.25	1.17
	LoRaOHRF	**0.91**	0.67	1.27	**0.23**	**1.17**
	LoRaPCR	0.90	**0.33**	1.45	0.25	1.23

sion task to $v_{t_3}, v_{t_4}, v_{t_5}$. First, we can transform the above problem into a GSMT problem, as illustrated in Figure 7(b). Furthermore, we can organize the special nodes into two connected components in Figure 7(c). The reduced GSMT problem is depicted in Figure 7(d).Consequently, the algorithm's complexity reduces from $O(n^c \cdot 2^m)$ to $O(n^c \cdot 2^{m'})$, where m' represents the number of connected components. In conclusion, through the Connected Components Reduction, we achieve an exponential speed improvement. The pseudocode for the Optimized Dynamic Programming Algorithm with Reduction is shown in Algorithm 3.

To further optimize the algorithm's performance, we apply pruning techniques for the cases where $m' = 1$ and $m' = 2$. The former is equivalent to finding the minimum spanning tree, while the latter corresponds to computing the shortest path between two connected components.

IV. EVALUATION

A. Methodology

We conduct extensive simulations using the CARLA simulator to evaluate the performance of LoRaPCR concerning various vehicle spatial distributions and different registration request scenarios. Specifically, we select three different maps in the CARLA simulator due to their distinct road characteristics, which may affect the spatial distribution of vehicles. We set the number of vehicles in the scene to be from 10 to 100 at an interval of 10 to test LoRaPCR's performance under different vehicle densities.

To generate vehicle registration requests, each vehicle selects up to REQ_{max} vehicles from the interest region (*i.e.* within a range of 100 meters) based on their preferences. We test the cases where REQ_{max} is set to 5, 10, 15, and 20, and we define three different preference as follows: 1)uniform:

refers to evenly assign the probability of being selected to all vehicles. 2)near: refers to assign a weight of $e^{(100-d)/20}$ to vehicles d meters away from the interest region and then normalized to obtain the probability. 3)far: refers to assign a weight of $e^{d/20}$ to vehicles d meters away from the interest region and then normalized to obtain the probability.

We collect point clouds from vehicle-mounted LiDARs at various distances in the CARLA simulator. A pre-trained GCL model is used for registration to fit the quality of single-hop point cloud registration (*i.e.* recall, RTE and RRE) across different distances and to analyze the error accumulation in multi-hop point cloud registration. For each PCR task, we estimate the registration success probability and error based on the distance between point cloud pairs using the registration quality estimation model fitted in the aforementioned process.

We define four metrics as follows:

- **Request Satisfaction Rate (RSR)**: refers to the ratio of the number of requests for which registration both path and transmission route are successfully found to the total number of requests.
- **Registration Cost Per Request (RCPR)**: refers to the average number of registrations performed for each registration request, indicating the average computational overhead.
- **Communication Cost Per Request (CCPR)**: refers to the average number of point cloud transmissions performed for each registration request, indicating the average communication overhead.
- **RTE** & **RRE**: refers to the average estimated upper bound of RTE and RRE for the registration results of each satisfiable registration request.

We compare our method with two candidate methods as follows:

- **Traditional direct registration methods(DRCT)**: refers to the approach of determining whether each request can be directly registered in one hop. If feasible, the registration and transmission are executed; otherwise, the request is rejected directly.
- **BFS-based registration methods (LoRaBFS)**: refers to using BFS to find a shortest path for each request without considering the reuse of duplicate edges, which

Fig. 8: Delivery ratio

Fig. 9: The CDF of m

is a straightforward method for implementing multi-hop registration.

B. Parameter configuration

1) Ablation Study: First, we evaluate the effectiveness of OHRF strategy and reduction on one-hop paths. We select the Town scenario in the CARLA simulator with $vehiclenumber = 80$ and $\text{REQ}_{max} = 10$. The maximum number of hops for registration paths is set to 7. We test the performance of LoRaBFS, LoRaOHRF, and LoRaPCR under three different registration request preferences.

Table I lists the performance of LoRaBFS, LoRaOHRF and LoRaPCR under three different registration request preferences. It can be seen that LoRaOHRF reduces RCPR to a certain extent, while improving CCPR, RTE, and RRE slightly. This is likely because determining the one-hop request paths first helps in finding better registration paths and transmission routes for hard requests. Furthermore, reduction on one-hop paths further decreases RCPR while slightly reducing the other metrics, which still remain within an acceptable range.

2) Impact of communication range: In the experiments, we assume that vehicles could engage in direct high-speed V2V communication with other vehicles within their communication range to support point cloud transmission tasks. The communication range of vehicles can influence the completion of point cloud transmission tasks. We define the delivery ratio as the ratio of the number of satisfied registration requests to the total number of registration requests that have feasible registration paths. We measure the delivery ratio of registration requests under different communication ranges, and the results are shown in Figure 8. When the communication range equals the interest region radius(*i.e.* 100 meters), all point cloud transmission tasks can be completed through one-hop transmission. As the communication range decreases, some point cloud transmission tasks may require multi-hop transmission or may become unreachable, leading to a decrease in the delivery ratio. It can be observed that the delivery ratio continuously increases as the communication distance expands. When the communication distance reaches 45m, the delivery ratio reaches 95%. Therefore, in the subsequent experiments, we set the communication distance to 45m.

3) Impact of the Connected Component Reduction: To evaluate the effectiveness of the Connected Component Reduction in the Transmission Task Scheduler, we conduct experiments under the same conditions as in the ablation study and record the cumulative distribution of the problem size. As shown in Figure 9, the Connected Component Reduction effectively
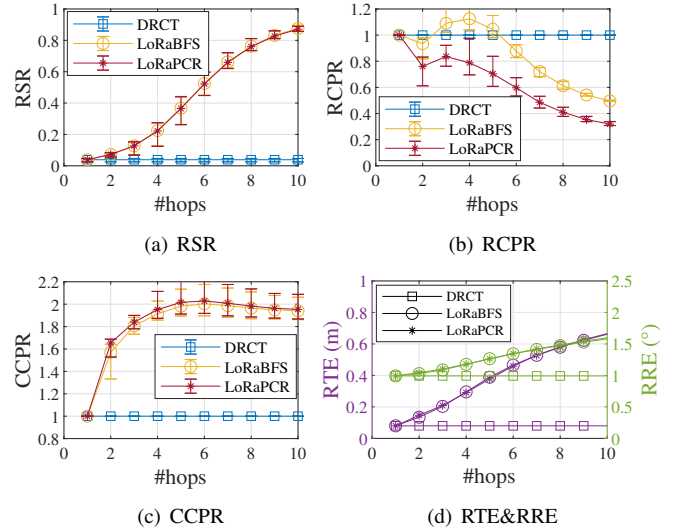

(a) RSR

(b) RCPR

(c) CCPR

(d) RTE&RRE

Fig. 10: Performance of DRCT, LoRaBFS and LoRaPCR with different maximum number of hops

(a) RSR

(b) RCPR

(c) CCPR

(d) RTE&RRE

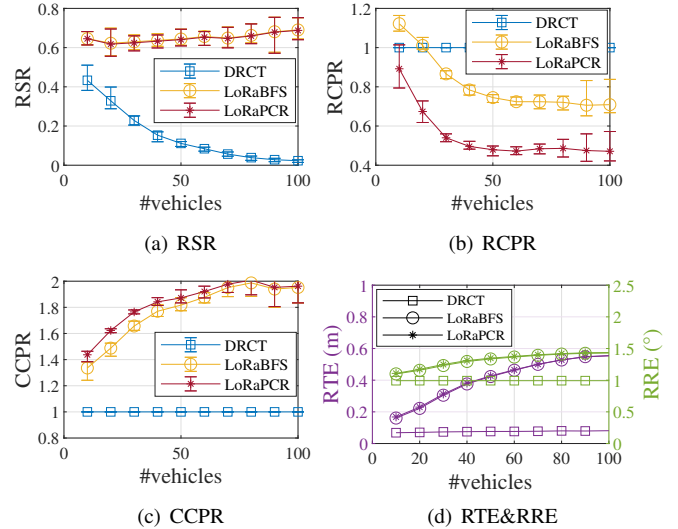Fig. 11: Performance of DRCT, LoRaBFS and LoRaPCR with different vehicle density

reduces the problem size, and in most of the cases(*i.e.* >80%), the problem size (the value of m) does not exceed 5, which is acceptable for the DP algorithm with a time complexity of $O(n^c 2^m)$. In most of the cases, the reduced algorithm can complete the computation within 1 ms, achieving a speedup of more than 1000 times compared to the non-reduced version.

4) Impact of maximum number of hops: We control the error of registration paths by limiting the maximum number of hops. We vary the maximum number of hops from 1 to 10 while keeping the remaining conditions consistent with the ablation study. We test the performance of DRCT, LoRaBFS and LoRaPCR under far preference.

Figure 10 shows that LoRaPCR demonstrates significant advantages in RCPR, while performing comparably to LoRaBFS in other metrics. Although DRCT demonstrates good per-
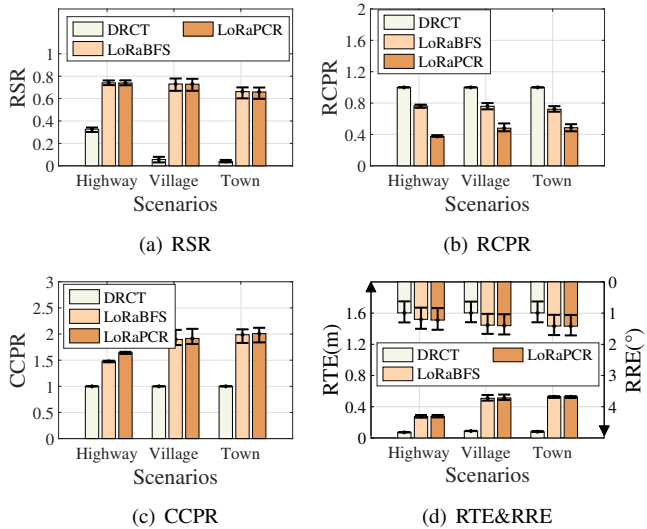
(a) RSR   (b) RCPR

(c) CCPR   (d) RTE&RRE

Fig. 12: Performance of DRCT, LoRaBFS and LoRaPCR in different scenarios

Fig. 13: Performance of DRCT, LoRaBFS and LoRaPCR in different REQ_{max}

formance in CCPR and RTE & RRE, its low RSR shows that it falls short of meeting the desired requirements. The results of the above experiments also validate our theoretical findings. When $HOP_{max} = 1$, both LoRaBFS and LoRaPCR degenerate into the DRCT algorithm. When $HOP_{max} = 2$, the system uses neighboring point cloud registration results for multi-hop registration, but it does not use previous time step results (which would require at least 3 hops). As a result, edge reuse starts to occur, leading to a decrease in RCPR. However, multi-hop registration paths introduce additional one-hop transmission tasks, resulting in an increase in CCPR. As HOP_{max} gradually increases, the system relaxes the requirements on the quality of registration paths, leading to an increase in RSR as more requests can be satisfied. Although more requests are fulfilled, the number of registration tasks in the system reaches saturation, meaning the total registration number and the additional transmission tasks do not increase. Consequently, RCPR increases with the growth of RSR, while CCPR remains stable (the impact of additional transmission tasks gradually becomes negligible in the overall transmission task).

From Figure 10(d), we can observe that when HOP_{max} is less than 8, the average RTE is below 0.6m, and the average RRE is below $1.5°$. Therefore, we set $HOP_{max} = 7$ in the following experiments.

C. Impact of Vehicle Density

To demonstrate the impact of different vehicle densities on the performance of LoRaPCR, we vary the number of vehicles from 10 to 100 at intervals of 10 and test the performance of LoRaPCR in the Town scenario with $REQ_{max} = 10$ under three different registration request preferences.

As shown in Figure 11, when the vehicle density is low, the advantages of multi-hop registration are not very prominent. However, as the vehicle density increases, the RCPR of multi-
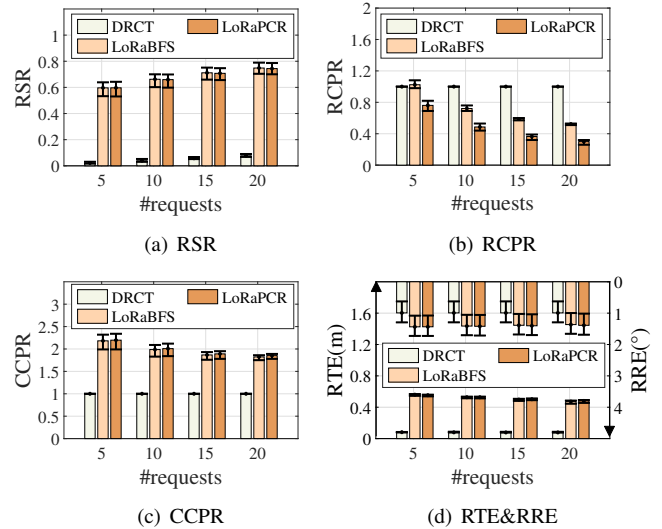
hop registration exhibits a significant improvement. It should be noted that when the vehicle density is very low, the number of vehicles within the interest region of a vehicle may not reach the threshold of REQ_{max}, invalidating the registration request preference and leading to relatively easy requests. As the vehicle density increases, the far preference begins to take effect, making the requests more challenging, resulting in an increase in CCPR, RTE and RRE. Meanwhile, RSR also shows a slight increase because of the increase of potential relay vehicles, but due to the aforementioned factors, the improvement is relatively small. Throughout these conditions, LoRaPCR consistently maintains a considerable advantage in RCPR, while its growth in CCPR remains within an acceptable range.

D. Impact of Scenarios

The distinct characteristics of different scenarios' roads can impact the spatial distribution of vehicles. We test the performance of LoRaPCR in three different scenarios while keeping the other conditions consistent with the ablation study. The results in Figure 12 show that multi-hop registration methods (LoRaBFS and LoRaPCR) can maintain a high RSR across various scenarios. LoRaPCR exhibits a significant advantage in RCPR compared to LoRaBFS, and the increased overhead in CCPR remains within an acceptable range. Both methods have similar registration errors.

E. Impact of maximum number of requests

The maximum number of requests for single vehicle affects the overall number of requests in the VANETs. We test the performance of LoRaPCR under different values (5, 10, 15, 20) of REQ_{max}, while keeping the other conditions consistent with the ablation study. Figure 13 shows that LoRaPCR maintains a significant advantage in RCPR compared to the other two methods while the disadvantage of CCPR remains

within an acceptable range. As REQ_{max} gradually decreases, the far preference becomes more effective, making the registration requests more challenging. Despite this, LoRaPCR still maintains a significant advantage in RCPR and performs roughly on par with LoRaBFS in other metrics.

V. DISCUSSION

A. Low Vehicle Density

When the vehicle density is low, it is possible that there may be no relay vehicle between two distant vehicles, making it hard to satisfy the registration request. In such cases, LoRaPCR attempts to utilize temporal point cloud data as a relay. If in a previous time instance the two vehicles are close enough to perform PCR or there is a relay vehicle in between, the current registration request can still be satisfied by multi-hop registration. In practical scenarios, if vehicles are very sparse and the distance between them is large, we can consider it relatively safe, and there may not be a strong demand for PCR.

B. Individual Communication Load Optimization

LoRaPCR is aiming to minimize the global communication load while overlooking the individual communication load of each vehicle. If a particular vehicle serves as a central node in the communication network, its communication load may become excessively high. To address this issue, we can introduce an algorithm to adjust the communication routes. In the MTG problem, the optimal transmission tree can have multiple solutions, which means that if a node's communication load is high, we can choose an alternative optimal transmission tree to avoid selecting that node while ensuring a constant global communication load remains unchanged.. If there is no other viable route for a specific transmission task, we can only use a "best-effort" strategy to satisfy the registration request to the best of the systems' communication capacity.

VI. RELATED WORK

We classify the existing relevant work into two categories: indoor PCR and outdoor PCR.

Indoor PCR. Indoor PCR can be further categorized based on the matched features into two groups: patch-based features and fully convolutional features. 3DMatch [13] extracts local features by employing 3D convolutions on local areas. PPF-Net [14] utilizes robust point-pair features extracted by PointNet for registration. PerfectMatch [15] adopts a voxelized smoothed density value (SDV) representation to obtain robust features. DIP [16] develops an effective method to register point clouds without initial alignment using distinctive 3D local deep descriptors. The latest advancements in SpinNet [17] and BUFFER [18] integrate cylindrical features that are equivalent under SO(2) rotations with backbones that are entirely convolutional. However, all these methods are local patch-based, which incurs significant computational overhead and cannot accomplish online registration of large-scale outdoor point clouds.

Outdoor PCR. Outdoor PCR can be further divided into two categories: pairwise V2V registration and multi-vehicle registration with infrastructure.

Pairwise V2V registration aims to optimize registration methods to improve registration performance. FCGF [19] introduces metric learning and utilizes sparse convolutions to accelerate computation. D3Feat [20] proposes a keypoint selection strategy to overcome inherent density variations. Predator [21] uses an overlap attention module to address the low overlap problem in PCR. APR [6] leverages an autoencoder to reconstruct a denser aggregated point cloud, allowing the encoder to extract features with rich local geometry information, thereby enhancing PCR accuracy. However, the aforementioned methods face difficulties in achieving PCR with distance exceeding 30m, which fails to meet the demands of autonomous driving.

Multi-vehicle registration with infrastructure methods enhances registration quality by constructing registration systems in VANETs. VI-Eye [22] leverages regular geometries in driving scenes to extract saliency points and efficiently achieve registration, enabling real-time PCR. However, the perception range of vehicles is still limited. VIPS [23] recognizes objects and utilizes their lean representations to construct a graph for matching, effectively extending the vehicle's perception range. However, the application scope of VIPS is limited due to the fact that vehicles do not share raw point cloud data. Both of the above methods require point cloud overlap between vehicles and infrastructure, which limits the coverage of infrastructure and leads to high deployment costs. EMP [5] utilizes Voronoi diagrams [24] and bandwidth considerations to partition regions for point cloud fusion. It can significantly enhance the perception range of vehicles. However, this system relies on infrastructure as an intermediary for point cloud transmission and depends on high-precision GPS to solve position relationships.

VII. CONCLUSION

In this work, we have developed an online long-range multi-vehicle PCR system called LoRaPCR. In LoRaPCR, vehicles can achieve long-range registration through multi-hop short-range high-quality registration. By collecting vehicle position information and registration history, the base station can assign efficient registration and transmission strategies for every registration request, reducing computation and communication overhead in VANETs. To the best of our knowledge, LoRaPCR is the first solution to achieve multi-vehicle point cloud long-range registration. We have implemented a prototype of LoRaPCR and conducted extensive simulations. The results demonstrate the effectiveness of LoRaPCR.

REFERENCES

[1] H. Zhu, J. Deng, Y. Zhang, J. Ji, Q. Mao, H. Li, and Y. Zhang, "VPFNet: Improving 3D Object Detection with Virtual Point based LiDAR and Stereo Data Fusion," *IEEE Transactions on Multimedia*, 2022.

[2] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander, "Joint 3D Proposal Generation and Object Detection from View Aggregation," in *Proceedings of IEEE IROS 2018*, pp. 1–8.

[3] C. Xu, B. Wu, Z. Wang, W. Zhan, P. Vajda, K. Keutzer, and M. Tomizuka, "SqueezeSegV3: Spatially-Adaptive Convolution for Efficient Point-Cloud Segmentation," in *Proceedings of ECCV 2020*, pp. 1–19.

[4] J. Shen, H. Zhu, Y. Cai, B. Zhai, X. Wang, S. Chang, H. Cai, and M. Guo, "mmV2V: Combating One-hop Multicasting in Millimeter-wave Vehicular Networks," in *Proceedins of IEEE ICDCS 2022*, pp. 735–742.

[5] X. Zhang, A. Zhang, J. Sun, X. Zhu, Y. E. Guo, F. Qian, and Z. M. Mao, "EMP: Edge-assisted Multi-vehicle Perception," in *Proceedings of ACM MobiCom 2021*, pp. 545–558.

[6] Q. Liu, Y. Zhou, H. Zhu, S. Chang, and M. Guo, "APR: Online Distant Point Cloud Registration Through Aggregated Point Cloud Reconstruction," in *Proceedings of IJCAI 2023*.

[7] X. Bai, Z. Luo, L. Zhou, H. Chen, L. Li, Z. Hu, H. Fu, and C.-L. Tai, "PointDSC: Robust Point Cloud Registration Using Deep Spatial Consistency," in *Proceedings of IEEE CVPR 2021*, pp. 15 859–15 869.

[8] G. D. Pais, S. Ramalingam, V. M. Govindu, J. C. Nascimento, R. Chellappa, and P. Miraldo, "3DRegNet: A Deep Neural Network for 3D Point Registration," in *Proceedings of IEEE CVPR 2020*, pp. 7193–7203.

[9] Q. Liu, H. Zhu, Y. Zhou, H. Li, S. Chang, and M. Guo, "Density-invariant Features for Distant Point Cloud Registration," in *Proceedings of IEEE ICCV 2023*.

[10] F. K. Hwang and D. S. Richards, "Steiner Tree Problems," *Networks*, vol. 22, no. 1, pp. 55–89, 1992.

[11] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An Open Urban Driving Simulator," in *Proceedings of CoRL 2017*, pp. 1–16.

[12] P. K. Agarwal and M. Sharir, "Efficient Algorithms for Geometric Optimization," *ACM Computing Surveys (CSUR)*, vol. 30, no. 4, pp. 412–458, 1998.

[13] A. Zeng, S. Song, M. Nießner, M. Fisher, J. Xiao, and T. Funkhouser, "3DMatch: Learning Local Geometric Descriptors from RGB-D Reconstructions," in *Proceedings of IEEE CVPR 2017*, pp. 1802–1811.

[14] H. Deng, T. Birdal, and S. Ilic, "PPFNet: Global Context Aware Local Features for Robust 3D Point Matching," in *Proceedings of IEEE CVPR 2018*, pp. 195–205.

[15] Z. Gojcic, C. Zhou, J. D. Wegner, and A. Wieser, "The Perfect Match: 3D Point Cloud Matching with Smoothed Densities," in *Proceedings of IEEE CVPR 2019*, pp. 5545–5554.

[16] F. Poiesi and D. Boscaini, "Distinctive 3D Local Deep Descriptors," in *Proceedings of IEEE ICPR 2020*, pp. 5720–5727.

[17] S. Ao, Q. Hu, B. Yang, A. Markham, and Y. Guo, "SpinNet: Learning a general surface descriptor for 3D point cloud registration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 753–11 762.

[18] S. Ao, Q. Hu, H. Wang, K. Xu, and Y. Guo, "BUFFER: Balancing accuracy, efficiency, and generalizability in point cloud registration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1255–1264.

[19] C. Choy, J. Park, and V. Koltun, "Fully Convolutional Geometric Features," in *Proceedings of IEEE CVPR 2019*, pp. 8958–8966.

[20] X. Bai, Z. Luo, L. Zhou, H. Fu, L. Quan, and C.-L. Tai, "D3Feat: Joint Learning of Dense Detection and Description of 3D Local Features," in *Proceedings of IEEE CVPR 2020*, pp. 6359–6367.

[21] S. Huang, Z. Gojcic, M. Usvyatsov, A. Wieser, and K. Schindler, "PREDATOR: Registration of 3D Point Clouds with Low Overlap," in *Proceedings of IEEE CVPR 2021*, pp. 4267–4276.

[22] Y. He, L. Ma, Z. Jiang, Y. Tang, and G. Xing, "VI-Eye: Semantic-based 3D Point Cloud Registration for Infrastructure-assisted Autonomous Driving," in *Proceedings of ACM MobiCom 2021*, pp. 573–586.

[23] S. Shi, J. Cui, Z. Jiang, Z. Yan, G. Xing, J. Niu, and Z. Ouyang, "VIPS: Real-Time Perception Fusion for Infrastructure-Assisted Autonomous Driving," in *Proceedings of ACM MobiCom 2022*, pp. 133–146.

[24] F. Aurenhammer, "Voronoi Diagrams—A Survey of a Fundamental Geometric Data Structure," *ACM Computing Surveys (CSUR)*, vol. 23, no. 3, pp. 345–405, 1991.