RadCloud: Real-Time High-Resolution Point Cloud Generation Using Low-Cost mmWave Radars for Aerial and Ground Vehicles

David Hunt, Shaocheng Luo, Amir Khazraei, Xiao Zhang, Spencer Hallyburton, Tingjun Chen, Miroslav Pajic
Department of Electrical and Computer Engineering, Duke University

ABSTRACT

We demonstrate RadCloud, a *real-time* framework for obtaining high-resolution lidar-like 2D point clouds from low-resolution millimeter-wave (mmWave) radar data on resource-constrained platforms commonly found on unmanned aerial and ground vehicles (UAVs and UGVs). Such point clouds can then be used for mapping key features of the environment, route planning and navigation, and other robotics tasks. RadCloud is specifically optimized for UAVs and UGVs by using a radar configuration with 1/4th the range resolution, using a model with 2.25× fewer parameters, and reducing total sensing time by a factor of 250×. The real-time ROS framework will be demonstrated on a UGV and UAV equipped with CPU-only compute platforms in diverse environments.

CCS CONCEPTS

• **Computer systems organization** → *Real-time systems*; *Sensors and actuators*; Robotics.

ACM Reference Format:
David Hunt, Shaocheng Luo, Amir Khazraei, Xiao Zhang, Spencer Hallyburton, Tingjun Chen, Miroslav Pajic. 2024. RadCloud: Real-Time High-Resolution Point Cloud Generation Using Low-Cost mmWave Radars for Aerial and Ground Vehicles. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24), November 18–22, 2024, Washington D.C., DC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3636534.3698849

1 INTRODUCTION

When it comes to indoor mapping and navigation, light detection and ranging (lidar) sensors like the commonly used VLP-16 Puck have been the gold standard thanks to their excellent angular and range resolutions of up to 0.1° and 2 mm, respectively [9]. However, such sensors consume up to 8 W and are relatively large, heavy (>800 g), and expensive (>$4,500) compared to other sensors, making them poorly suited for cheap and small vehicles like unmanned aerial and ground vehicles (UAVs and UGVs, respectively) [1, 9].

On the other hand, millimeter-wave (mmWave) radio detection and ranging (radar) sensors like the popular IWR1443 from Texas Instruments only consumes 2 W and are significantly smaller, lighter (<250 g), and cheaper (<$400) than lidar sensors while achieving range resolutions of up to 4 cm due to their large signal bandwidth of 4 GHz[7, 10]. While these sensors are common in automotive applications like [5, 16], they are also ideal for UGV, UAV, and handheld applications [11, 15]. However, these sensors have poor angular resolution (e.g., only 30° for the TI-IWR1443 [13]).

In this work, we demonstrate the RadCloud framework [6], which can be used to generate high-resolution lidar-like 2D point clouds from low-resolution mmWave radar data in real-time, even on resource-constrained UAVs and UGVs. To achieve this, RadCloud builds upon previous work from [11, 12] which utilized a modified U-Net architecture [14] to convert radar data into lidar-like point clouds, but wasn't optimized for resource-constrained platforms which don't have a GPU and can only process raw radar data at relatively low rates. To overcome these challenges, RadCloud utilizes radar data with 1/4th the maximum possible range resolution and a model with 2.25× fewer parameters than the model from [11, 12], enabling real-time operation on CPU-only platforms such as the Intel NUC. Additionally,

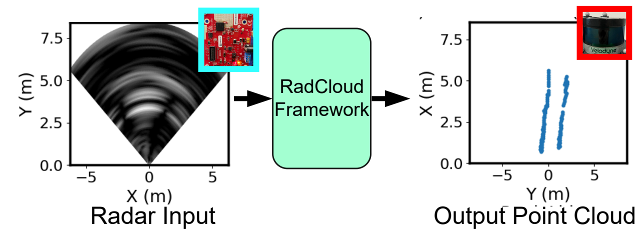

Figure 1: RadCloud [6] utilizes a modified U-Net architecture to convert low resolution radar range azimuth responses into high-resolution 2D point clouds

ACM ISBN 979-8-4007-0489-5/24/11
https://doi.org/10.1145/3636534.3698849

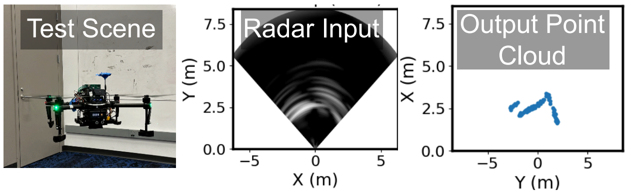

Figure 2: Example output from RadCloud [6] framework running in real-time on a UAV

we improve RadCloud's resiliency to rapid translations and turns by utilizing only data from a single radar frame for each generated point cloud. Even with these optimizations, we showed that RadCloud is still able to capture the key features of the environment with 90% of predicted point cloud points being <40 cm from the nearest ground truth lidar points. Additional resources, including code and datasets are available at [2].

2 OVERVIEW OF RADCLOUD

Fig. 1 presents an overview of the RadCloud [6] framework with the following key elements.

mmWave radar sensing: As shown in Fig. 1, the RadCloud framework takes in a series of radar range azimuth responses which represent the power of the reflected radar signals at different ranges and angles. Here, the radar field of view (FoV) is restricted to $\pm 50°$. To capture the raw data from the TI-IWR1443 mmWave radar we utilize the TI-DCA1000 data capture card. We configured our radar to operate at 1/4th the maximum possible resolution to accommodate for the limited sampling rate. The final radar configuration achieved a range resolution of 13.3 cm, maximum range of 8.56 m, and angular resolution of 28.6°.

RadCloud model: RadCloud utilizes a modified U-Net architecture with 7.7M parameters to obtain high-resolution 2D point clouds from low-resolution radar data [14], representing a 2.25× reduction in the model size compared to [11, 12]. For its input, RadCloud utilizes 40 radar "chirps" recorded during a single frame. This results in a 250 × reduction in sensing time compared to the previous work, significantly improving resiliency to rapid movements. At the output, RadCloud generates a 2D high-resolution point cloud from the input radar data.

Real-time ROS implementation: Finally, the entire RadCloud framework is implemented as a ROS package which we have made open-source to facilitate research in this area by the broader community. The ROS package includes nodes for streaming raw radar data from the TI-IWR1443 and TI-DCA1000, processing the radar data in real-time, and generating 2D point clouds using a pre-trained RadCloud model. The entire framework works on resource-constrained platforms such as the NUC7i5BNH (a CPU-only platform) where we achieved average frame rates above 15 frames per second when running the entire pipeline.

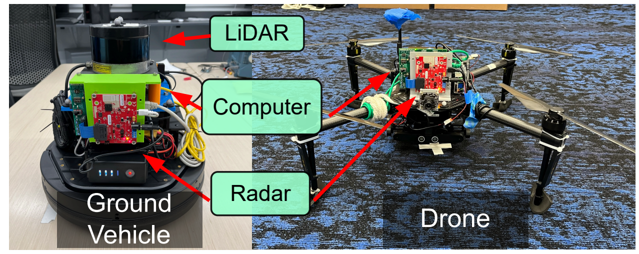

Figure 3: UAV and UGV demonstration platforms.

3 DEMONSTRATION

In this demonstration, we showcase the real-time RadCloud implementation running on both an unmanned ground vehicle and unmanned aerial vehicle as shown in Fig. 3. Both platforms will run the full RadCloud ROS stack at 10 frames per second, and monitors will display the input radar data, computed point cloud, and ground truth data from a lidar sensor using Rviz. Additionally, a laptop will be used to connect to and control both vehicles over a private network connection. During the demonstration, attendees can manually drive the UGV around in the environment to see how the RadCloud model performs in realistic scenarios. If the venue allows UAVs to be flown inside, we will manually fly the UAV to demonstrate performance on airborne platforms. We will bring all necessary materials including monitors, environmental obstacles, routers, and other compute hardware. Additional demo videos are featured on [3].

Hardware: The Kobuki UGV and DJI Matrice 100 UAV platforms shown in Fig. 3 will be utilized to demonstrate the capability of the RadCloud framework to generate point clouds in real-time [4, 8]. Both platforms are equipped with the TI TI-IWR1443 radar, TI-DCA1000 and an Intel NUC running the entire framework in real-time. Additionally, the Kobuki UGV is equipped with a lidar sensor to provide ground truth data for comparison.

Environment and setup: A 5 m × 5 m environment should be sufficient space for the demonstration. While we will bring several moveable dividers to reconfigure the demonstration environment for different scenes, it would be helpful for the demonstration environment to be located at the corners of the conference room with access to a power outlet.

ACKNOWLEDGMENTS

This work is sponsored in part by the ONR N00014-23-1-2206 and AFOSR FA9550-19-1-0169 awards, NSF under NAIAD Award 2332744, CNS-2211944, and AST-2232458 grants, and the National AI Institute for Edge Computing Leveraging Next Generation Wireless Networks, Grant CNS-2112562.

REFERENCES

[1] [n. d.]. Puck. https://store.clearpathrobotics.com/products/puck

[2] 2024. RadCloud. https://sites.google.com/view/radcloudduke

[3] 2024. RadCloud Case Studies. https://sites.google.com/view/radcloudduke/case-studies

[4] DJI. 2023. Matrice 100 Specs. https://www.dji.com/matrice100

[5] David Hunt, Kristen Angell, Z. Qi, Tingjun Chen, and Miroslav Pajic. 2024. MadRadar A Black-Box Physical Layer Attack Framework on mmWave Automotive FMCW Radars. In *The 2024 Network and Distributed System Security Symp. (NDSS)*. Can Diego, CA.

[6] David Hunt, Shaocheng Luo, Amir Khazraei, Xiao Zhang, Spencer Hallyburton, Tingjun Chen, and Miroslav Pajic. [n. d.]. RadCloud: Real-Time High-Resolution Point Cloud Generation Using Low-Cost Radars for Aerial and Ground Vehicles. In *2024 IEEE Int. Conf. on Robot. and Automat. (ICRA'24)*.

[7] Texas Instruments. 2018. IWR1443 Single-Chip 76- to 81GHz mmWave Sensor. https://www.ti.com/lit/gpn/iwr1443

[8] Kobuki. 2016. Kobuki-UserGuide.pdf. http://file.ncnynl.com/ros/Kobuki-UserGuide.pdf

[9] Velodyne LiDAR. 2018. VLP-16 User Manual. https://velodynelidar.com/wp-content/uploads/2019/12/63-9243-Rev-E-VLP-16-User-Manual.pdf

[10] Mouser. [n. d.]. IWR1443Boost. https://www.mouser.com/ProductDetail/Texas-Instruments/IWR1443BOOST?qs=5aG0NVq1C4wT7gyvvDbMRw%3D%3D

[11] Akarsh Prabhakara, Tao Jin, Arnav Das, Gantavya Bhatt, Lilly Kumari, Elahe Soltanaghai, Jeff Bilmes, Swarun Kumar, and Anthony Rowe. 2023. High Resolution Point Clouds from mmWave Radar. In *2023 IEEE Int. Conf. on Robot. and Automat. (ICRA'23)*. IEEE, London, United Kingdom, 4135–4142. https://doi.org/10.1109/ICRA48891.2023.10161429

[12] Akarsh Prabhakara, Tao Jin, Arnav Das, Gantavya Bhatt, Lilly Kumari, Elahe Soltanaghai, Jeff Bilmes, Swarun Kumar, and Anthony Rowe. 2023. RadarHD: Demonstrating Lidar-like Point Clouds from mmWave Radar. In *Proc. of the 29th Annu. Int. Conf. on Mobile Comput. and Netw. (MobiCom'23)*. 1–3.

[13] Sandeep Rao. [n. d.]. Introduction to mmwave Sensing: FMCW Radars. https://training.ti.com/sites/default/files/docs/mmwaveSensing-FMCW-offlineviewing_0.pdf

[14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. http://arxiv.org/abs/1505.04597 arXiv:1505.04597 [cs].

[15] Emerson Sie, Zikun Liu, and Deepak Vasisht. 2023. BatMobility: Towards Flying Without Seeing for Autonomous Drones. In *Proc. of the 29th Annu. Int. Conf. on Mobile Comput. and Netw. (MobiCom'23)* (Madrid, Spain) *(ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 31, 16 pages. https://doi.org/10.1145/3570361.3592532

[16] Elahe Soltanaghaei, Akarsh Prabhakara, Artur Balanuta, Matthew Anderson, Jan M. Rabaey, Swarun Kumar, and Anthony Rowe. 2021. Millimetro: mmWave retro-reflective tags for accurate, long range localization. In *Proc. of the 27th Annu. Int. Conf. on Mobile Comput. and Netw. (MobiCom'21* (New Orleans, Louisiana) *(MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 69–82. https://doi.org/10.1145/3447993.3448627

Received 7 August 2024; revised 28 September 2024; accepted 1 October 2024