2025 EMBEDDED VISION SUMMIT

# Optimizing Real-Time SLAM Performance for Autonomous Robots with GPU Acceleration

Naitik Nakrani

Ph.D., Solution Architect Manager

Einfochips, India

✉ robotics@einfochips.com

eInfochips
An Arrow Company

# Optimizing real-time SLAM performance for autonomous robots with GPU acceleration

## Outline

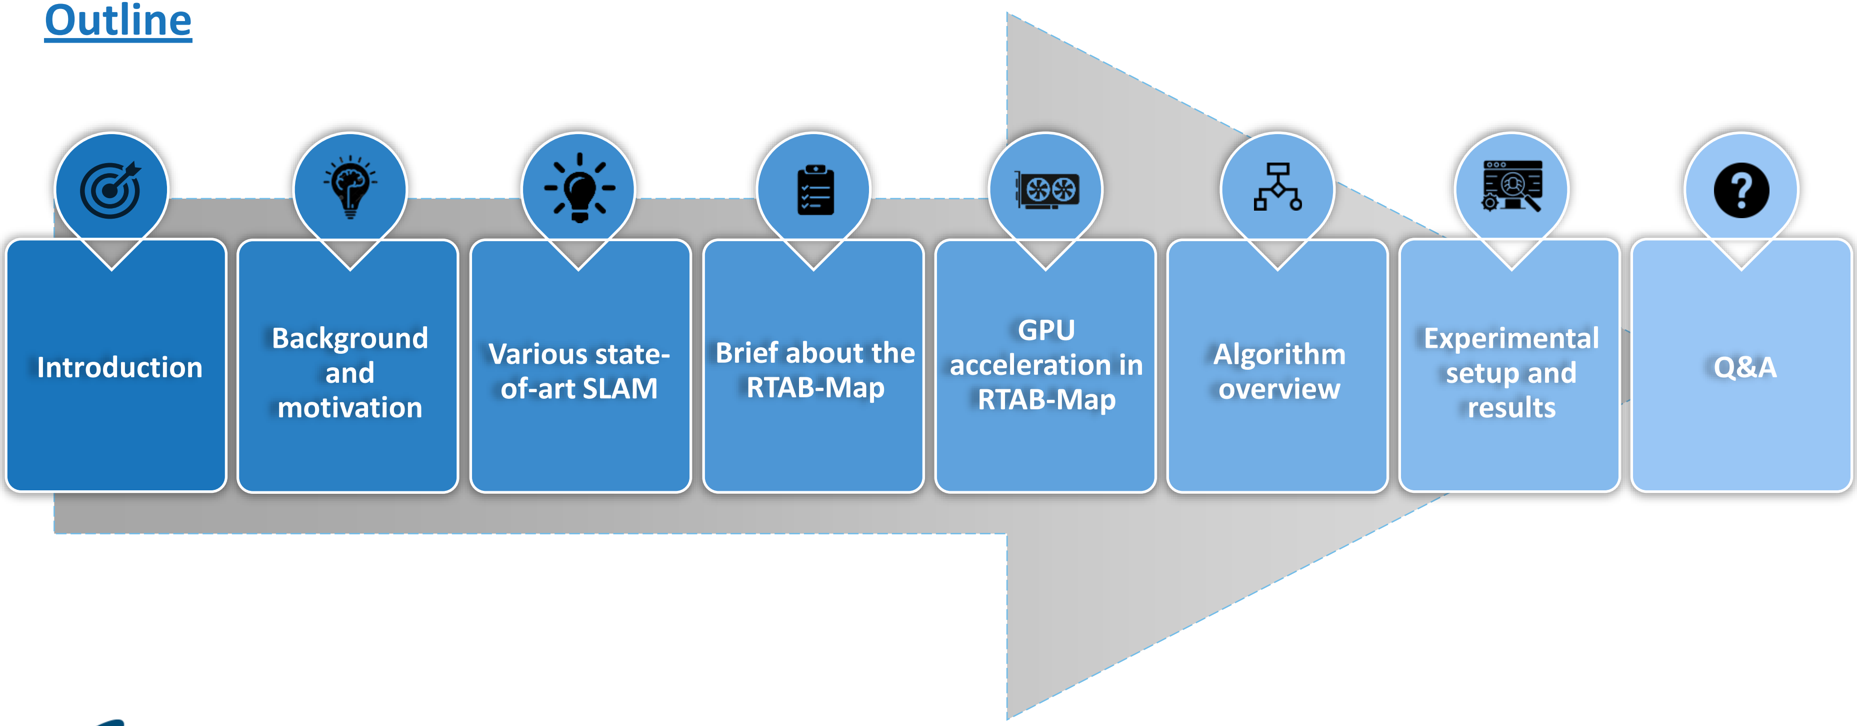

Introduction

Background and motivation

Various state-of-art SLAM

Brief about the RTAB-Map

GPU acceleration in RTAB-Map

Algorithm overview

Experimental setup and results

Q&A

# Introduction

SLAM is simultaneous localization and mapping

Key components:

- Localization: Estimating the robot's position and orientation within the environment.
- Mapping: Creating a spatial map of the surroundings using sensors (e.g., LiDAR, cameras).

Applications:

- Autonomous vehicles
- Robotics (e.g., vacuum cleaners, drones)
- AR/VR systems

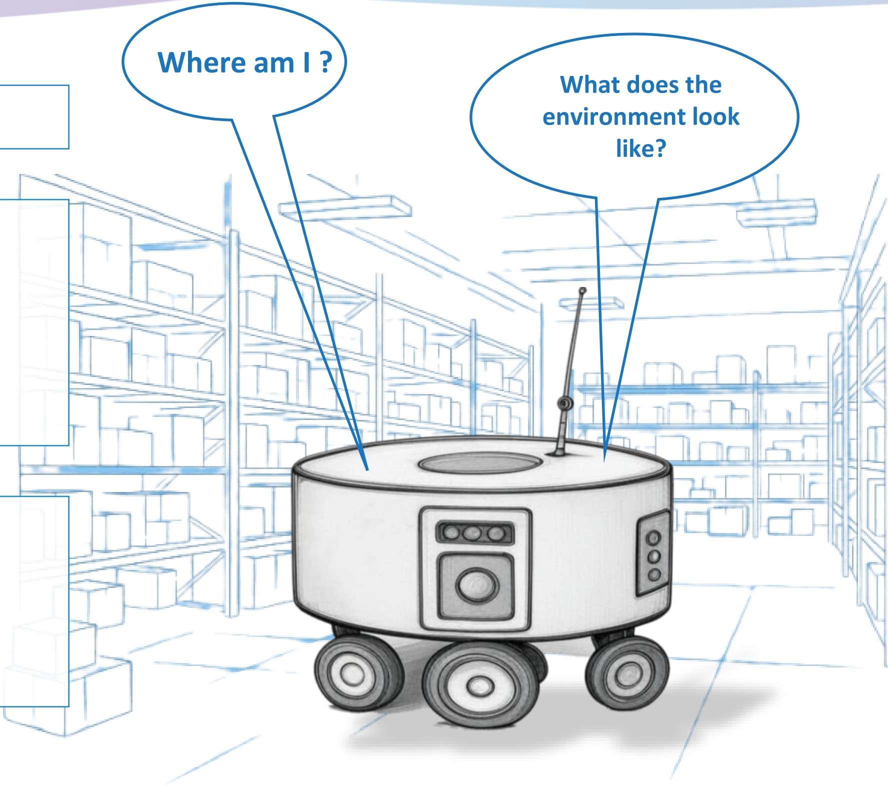

# Background and motivation

**There are several challenges that SLAM faces in real time:**

- Large scale and long-term
- **Processing time for each frame**
- Odometry drift
- Environmental variations
- Kidnapped robot problem
- No navigation support
- Dynamic obstacle

**Objective: To reduce the processing time by leveraging hardware acceleration**

# Various state-of-the-art SLAM

| | Input sources | | | | | | Outputs | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | Camera | | | Lidar | | Odom | Pose | Occupancy | | Point Cloud |
| | Stereo | RGB-D | Multi | 2D | 3D | | | 2D | 3D | |
| GMapping | | | | ✓ | | ✓ | ✓ | ✓ | | |
| Hector SLAM | | | | ✓ | | | ✓ | ✓ | | |
| ETHZASL-ICP | | | | ✓ | ✓ | ✓ | ✓ | ✓ | | Dense |
| Karto SLAM | | | | ✓ | | ✓ | ✓ | ✓ | | |
| Cartographer | | | | ✓ | ✓ | ✓ | ✓ | ✓ | | Dense |
| BLAM | | | | | ✓ | | ✓ | | | Dense |
| VINS-Mono | | | | | | | ✓ | | | |
| ORB-SLAM2 | ✓ | ✓ | | | | | | | | |
| S-PTAM | ✓ | | | | | | ✓ | | | Sparse |
| DVO-SLAM | | ✓ | | | | | ✓ | | | |
| MCPTAM | ✓ | | ✓ | | | | ✓ | | | Sparse |
| RGBDSLAMv2 | | ✓ | | | | ✓ | ✓ | | ✓ | Dense |
| RTAB-Map | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | **Dense** |

https://arxiv.org/abs/2403.06341

# Brief about RTAB - Map

RTAB-Map is an open-source framework for real-time Simultaneous Localization and Mapping (SLAM) with loop

**It has features like**

- Multi-sensor integration
- Appearance-based approach
- Robust localization
- 2D/3D occupancy map generation
- Multi-session mapping (kidnapped robot problem)
- ROS compatible

This talk focuses on RTAB-Map with pointcloud data from ToF as pointcloud is computationally heavy data.

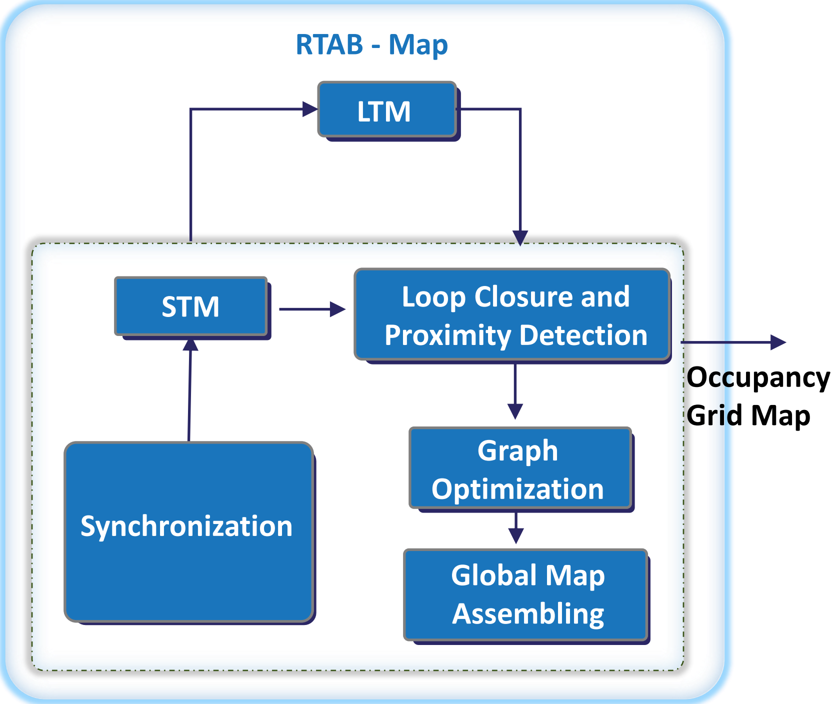

https://link.springer.com/chapter/10.1007/978-3-030-75472-3_7

# GPU acceleration in RTAB - Map

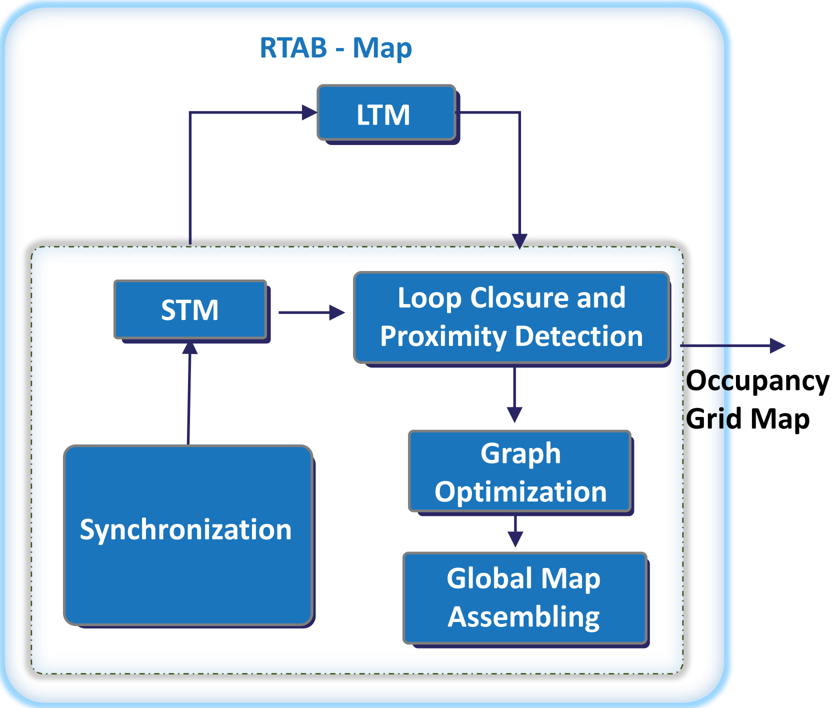

# GPU acceleration in RTAB - Map

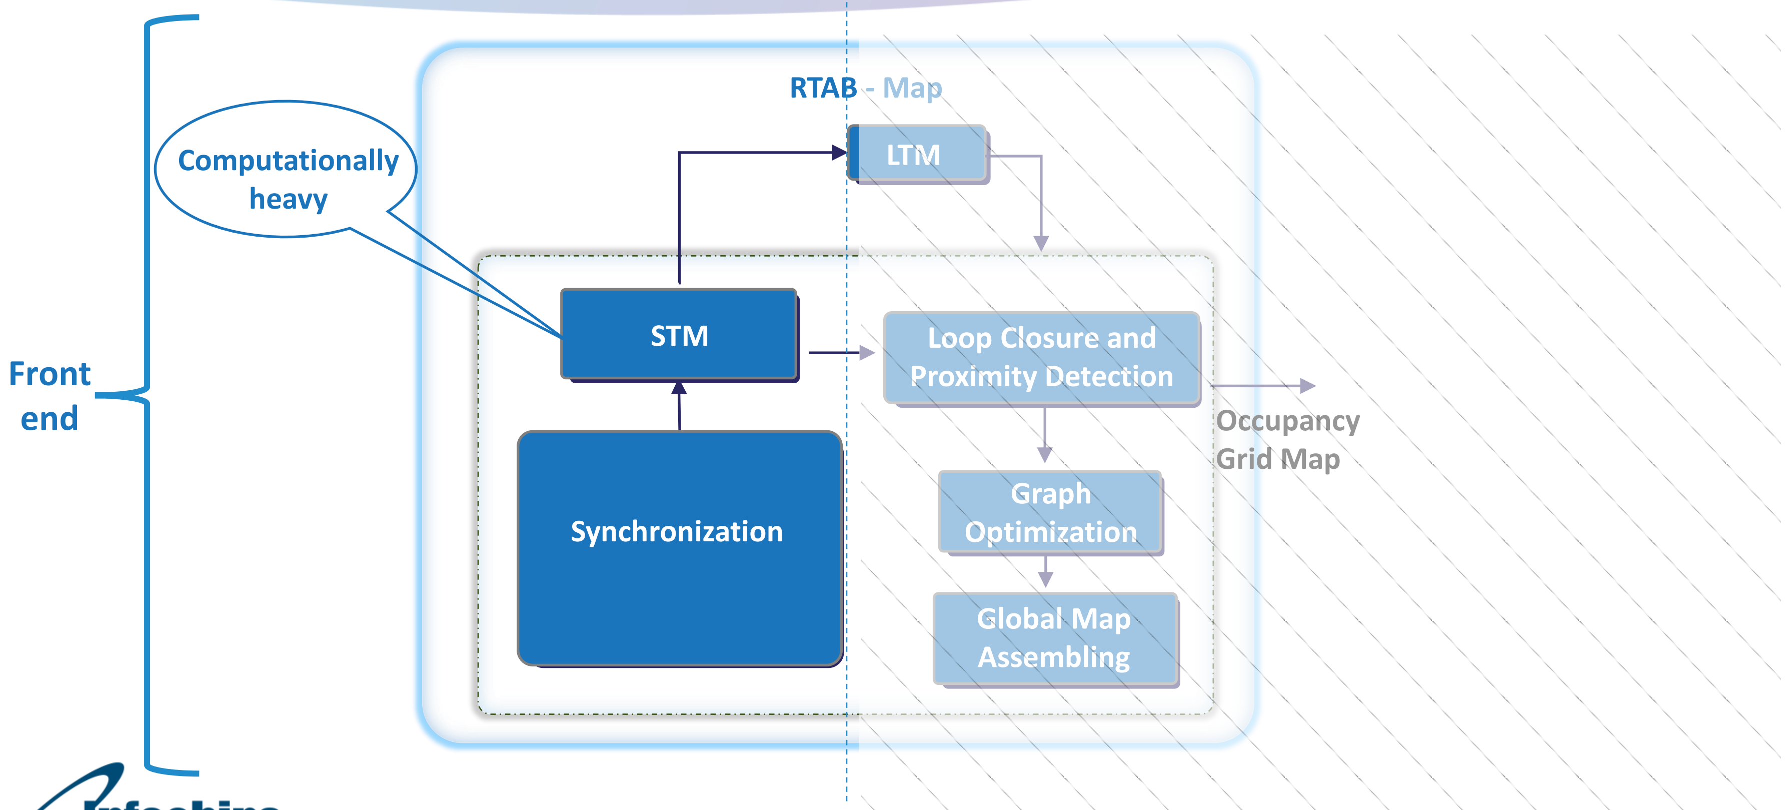

# GPU acceleration in RTAB - Map

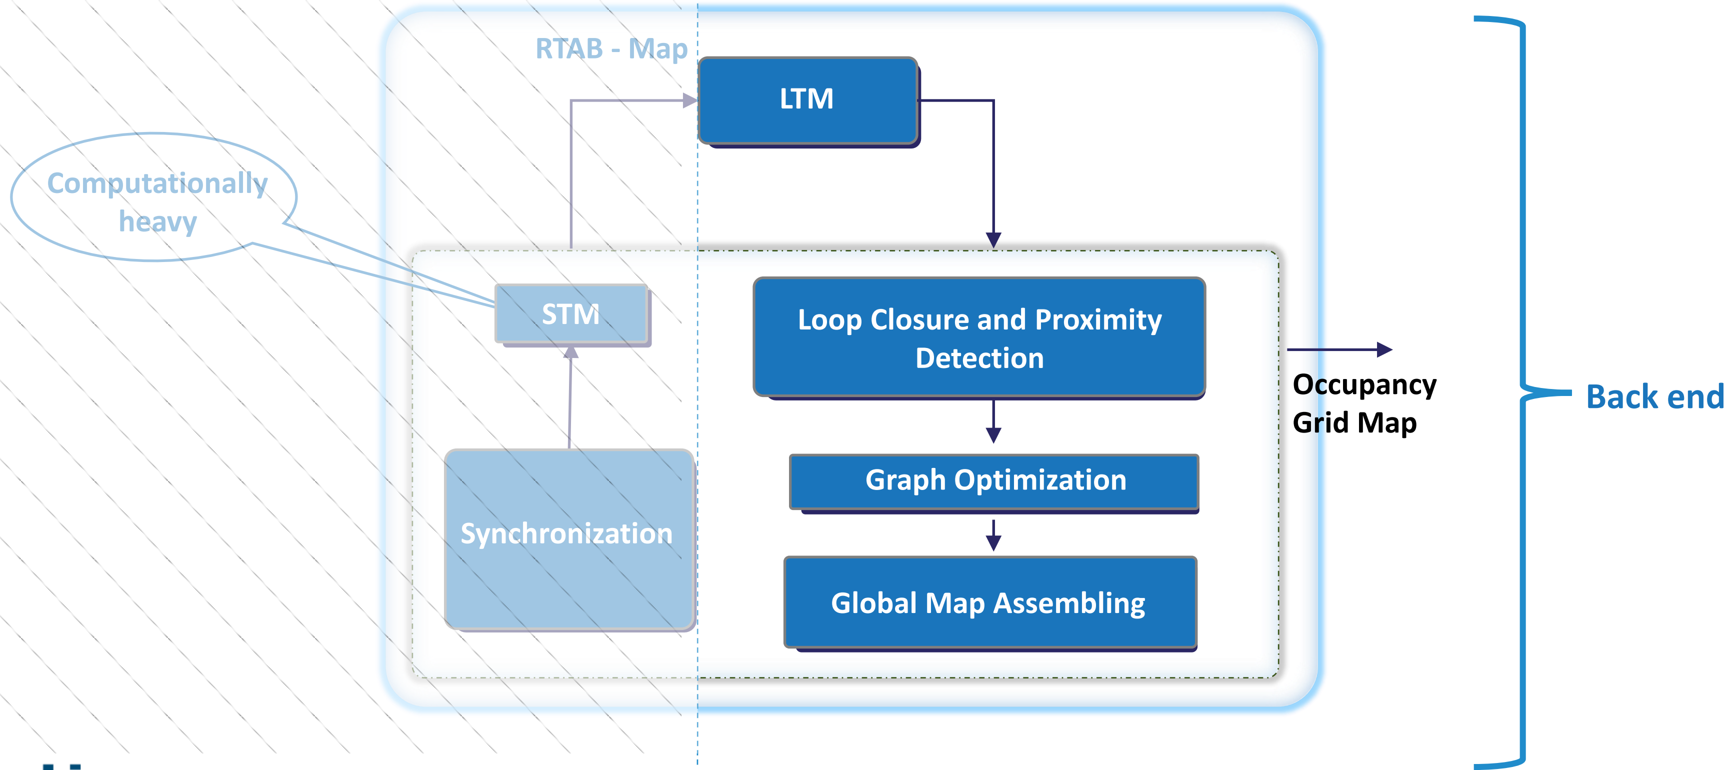

# Algorithm overview

# What is STM ?

STM is RTAB-Map block that performs data segmentation and generates occupancy map

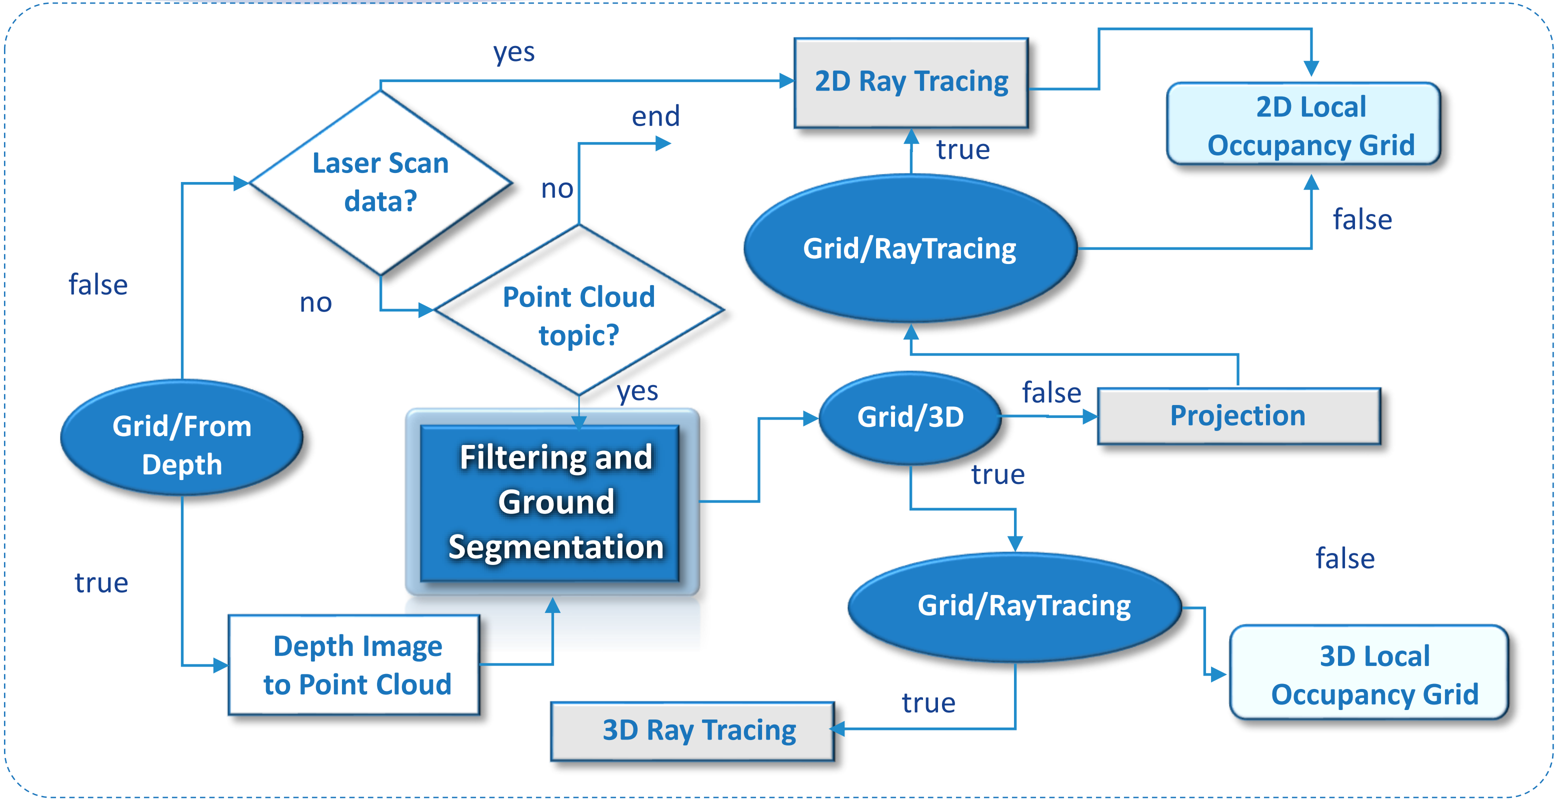

https://arxiv.org/abs/2403.06341

# Inside of filtering and ground segmentation

Segmentation is the process of dividing a pointcloud into smaller, meaningful segments or clusters.

RTAB-Map uses point cloud library (PCL) to segment a frame of pointcloud using Euclidean cluster extraction method.

Executes on CPU

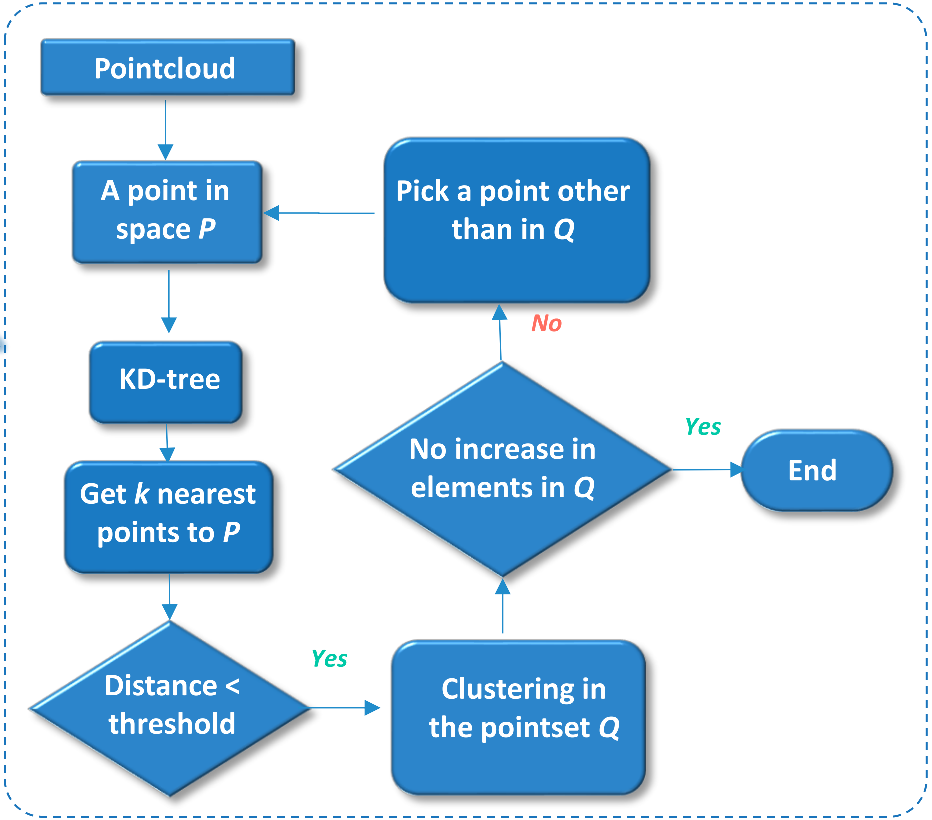

Pointcloud
A point in space *P*
KD-tree
Get *k* nearest points to *P*
Distance < threshold
Yes
Clustering in the pointset *Q*
No increase in elements in *Q*
No
Pick a point other than in *Q*
Yes
End

https://arxiv.org/abs/2403.06341

# Example of a Pointcloud segmentation

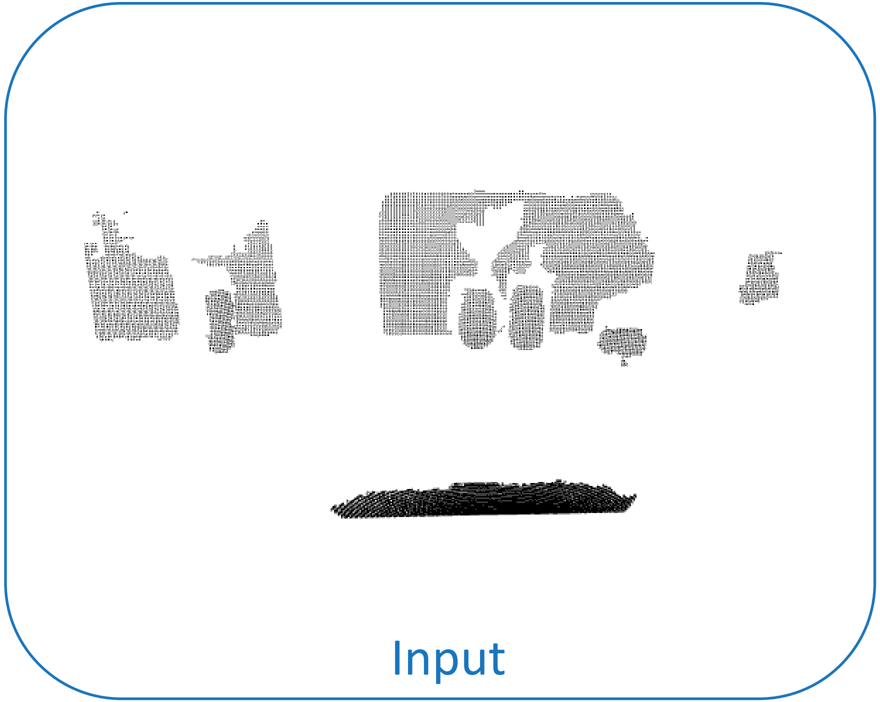

Input

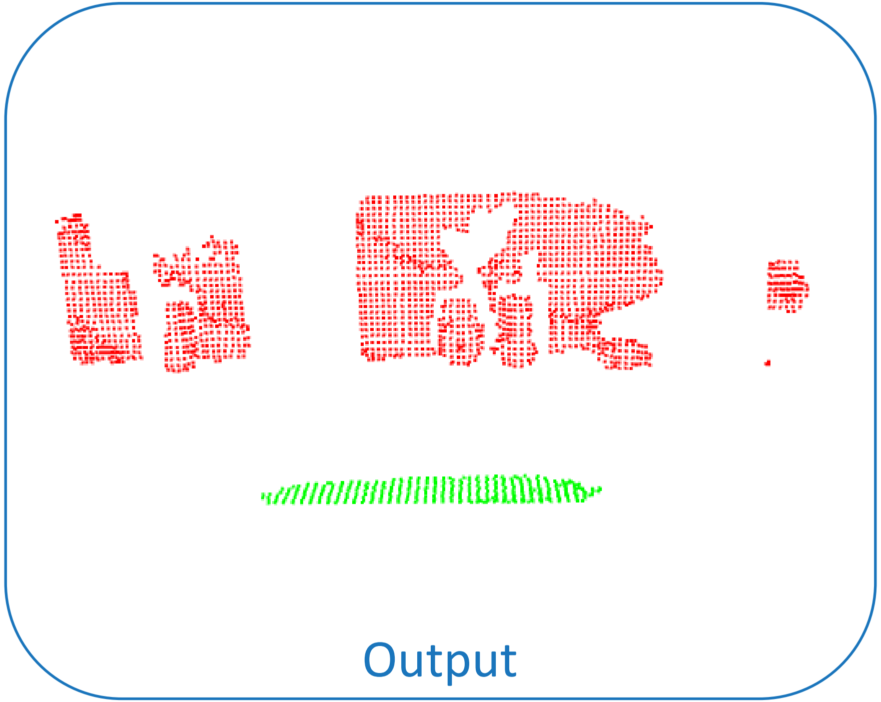

Output

© 2025 Einfochips (An Arrow Company)

# Proposed approach

KD-tree is efficient data structure for organizing point cloud but its computationally heavy.

A GPU based **Octree** was implemented

This approach allows us to transfer the computationally heavy pointcloud segmentation of each data frame from CPU to GPU.

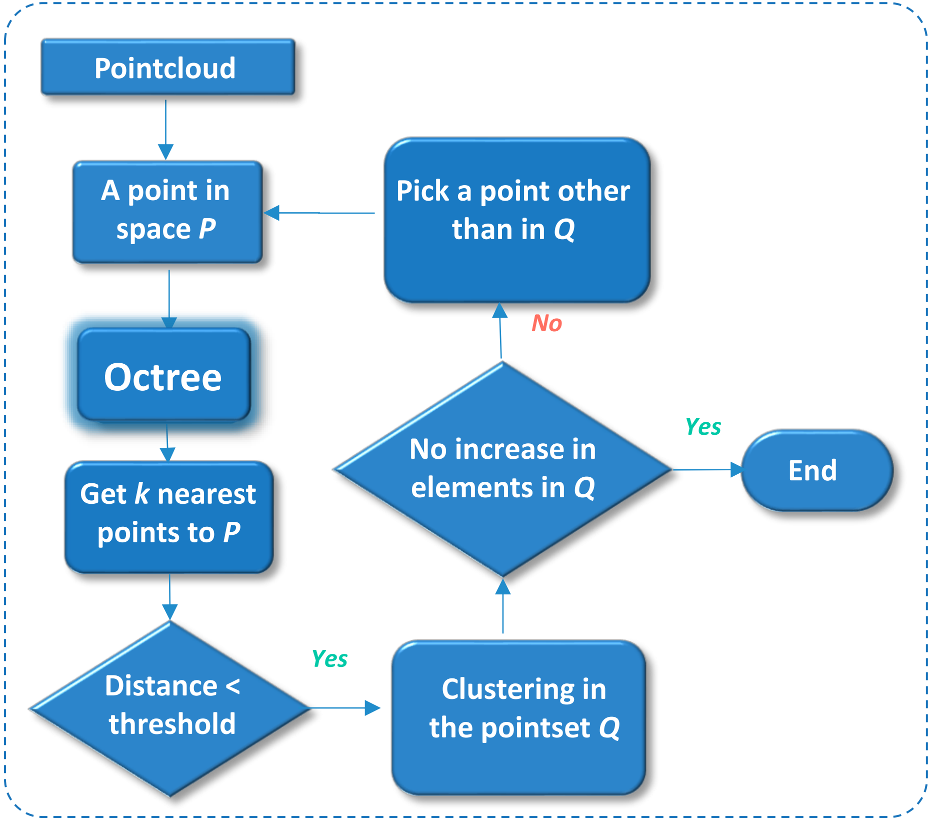
Pointcloud
A point in space *P*
Octree
Get *k* nearest points to *P*
Distance < threshold
Yes
Clustering in the pointset *Q*
No increase in elements in *Q*
Yes
End
No
Pick a point other than in *Q*

# Experimental setup and results

# Experimental setup

## Dataset

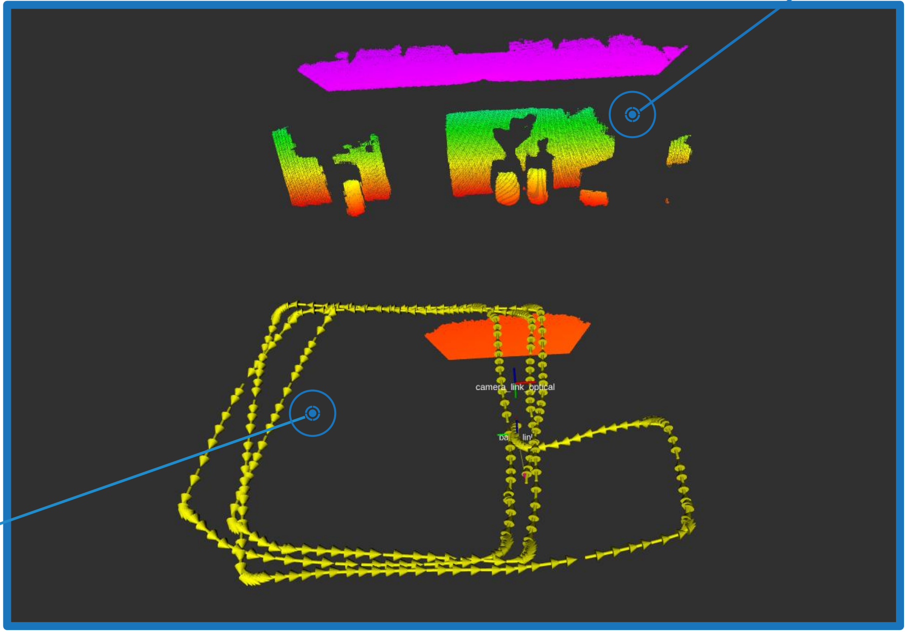

ToF Pointcloud

EKF based Odometry

ROS2 Humble,

Jetpack 5.1.2,

CUDA 11.4,

Containerized docker

# A ROS2 based RTAB-Map pipeline for evaluation

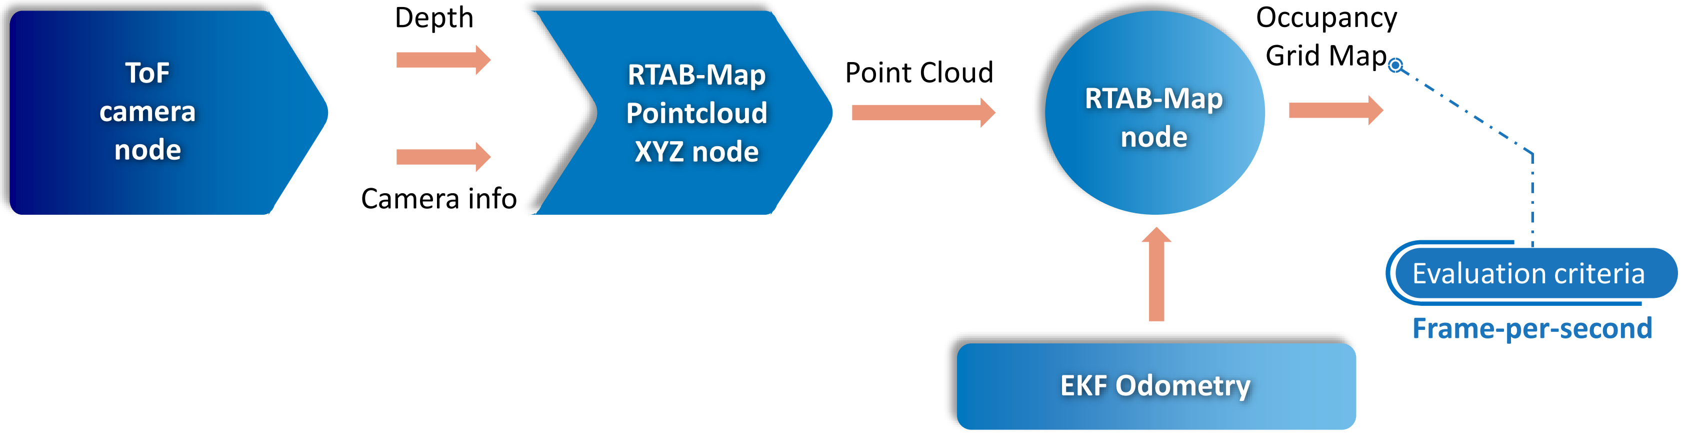

# Analysis and results

| Platform | Method | Segmentation Execution Time (ms) | Total RTAB-Map node execution time (ms) |
|---|---|---|---|
| Orin AGX | CPU | 242 | 398 |
| | GPU | 56 | 182 |
| | GPU-DR | 26 | 102 |
| Orin Nano | CPU | 577 | 667 |
| | GPU | 367 | 454 |
| | GPU-DR | 113 | 295 |

- Input pointcloud size is 512 x 512 at 18 FPS.
- For **GPU-Data reduction**, input pointcloud frame is cropped 20% and 10% from top and bottom, and decimated by factor 2

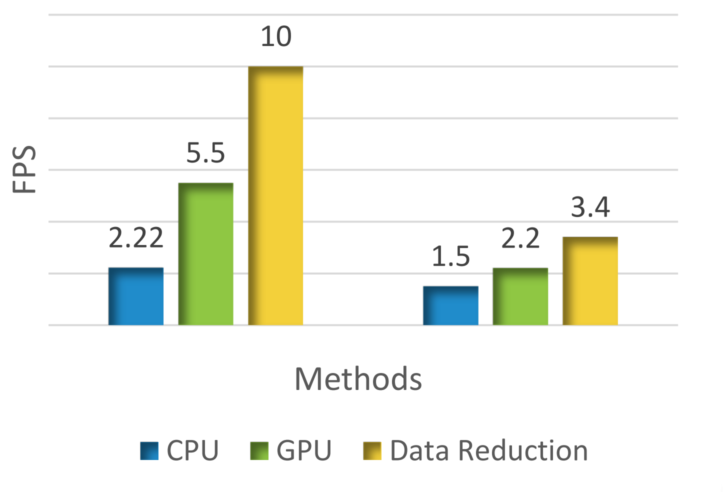

| | Max CPU freq | Max GPU freq |
|---|---|---|
| Orin AGX | 2.2 GHz | 1.3 GHz |
| Orin Nano | 1.5 GHz | 1.02 GHz |

# Analysis and results - map comparison vs FPS

| | | | |
|---|---|---|---|
| Orin AGX | 2.22<br>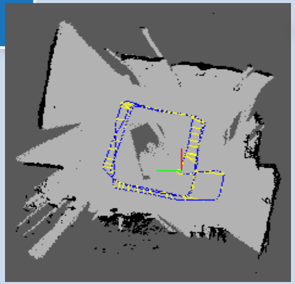 | 5.5<br>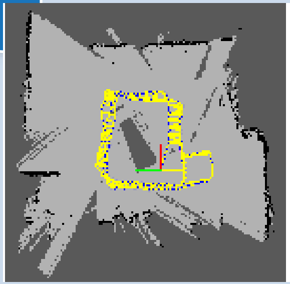 | 10<br>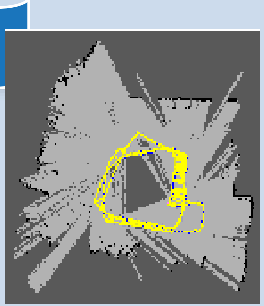 |
| Orin Nano | 1.5<br>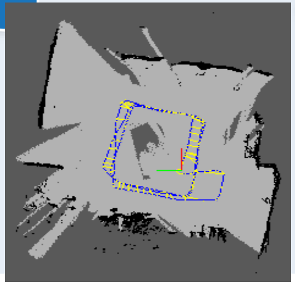 | 2.2<br>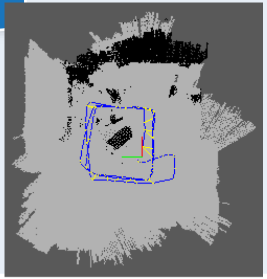 | 3.4<br>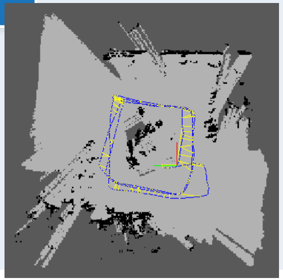 |

# Conclusion

1 For long-term SLAM, ideally SLAM should be performed in real-time, but practically it is challenging.

2 We managed to reduce the processing time as low as possible for each pointcloud frame so that overall map generating rate for RTAB-Map could be increased on a given edge platform.

3 Results show that by using octree and data reduction, the overall map generating rate (FPS) was increased by 4.5x and 2.25x for NVIDIA's Orin AGX and Orin Nano platforms, respectively.

# Resources

- ROS2 Humble ROS 2 Documentation — ROS 2 Documentation: Humble documentation
- RTAB-Map GitHub - https://github.com/introlab/rtabmap
- PCL - GitHub - PointCloudLibrary/pcl: Point Cloud Library (PCL)
- https://arxiv.org/abs/2403.06341
- https://link.springer.com/chapter/10.1007/978-3-030-75472-3_7
- https://www.researchgate.net/publication/377165321_Extraction_of_Feature_Information_from_Point_Cloud_of_Large_Volume_Steel_Truss_Members

## 2025 Embedded Vision Summit

## EInfochips

## Booth # 802

## robotics@einfochips.com

At the same time, you can also attend our talk sharing valuable insights on

“Efficiently Registering Depth and RGB Images” - Wednesday, May 21: 4:50 pm - 5:20 pm

# Q & A