

AFS: Accurate, Fast, and Scalable Cooperative Perception

Fawad Ahmad, Rochester Institute of Technology

Problem Statement:

- On-board vehicle 3D sensors are prone to occlusions (Fig 1).
- Cooperative perception shares 3D data among agents.
- Core component is 3D spatial alignment.
- 3D maps/GPS for alignment can be erroneous (Fig 2).

Challenges:

- Cooperative perception must be accurate, fast and scalable but there is a traditional tradeoff between accuracy and latency.
- Indirectly aligning point clouds is fast but inaccurate whereas directly aligning them accurate but can incur latency.

Contributions:

- We propose using multi-tiered alignment in which groups of vehicles align to dynamic anchors (vehicles or infrastructure-sensors) (Fig 3).
- To reduce latency, we find overlapping regions in point clouds and compute alignment from those regions.
- To reduce network bandwidth, we use spatial reasoning algorithms to deliver only relevant content to vehicles.

Results:

- With 40+ vehicles, our proposed work has a mean alignment accuracy of 2 cm at an end-to-end latency less than 20 ms (Fig 4).

Broader Impacts:

- Cooperative perception has the potential to significantly reduce accidents at traffic intersections.
- Cooperative perception can also be applied to other use cases like disaster relief.
- PhD students presented their ongoing work at a seminar class that I teach on autonomous driving.
- Half the students from the class are pursuing independent research with me this semester.




Fig 1: Cooperative perception can overcome occlusions.

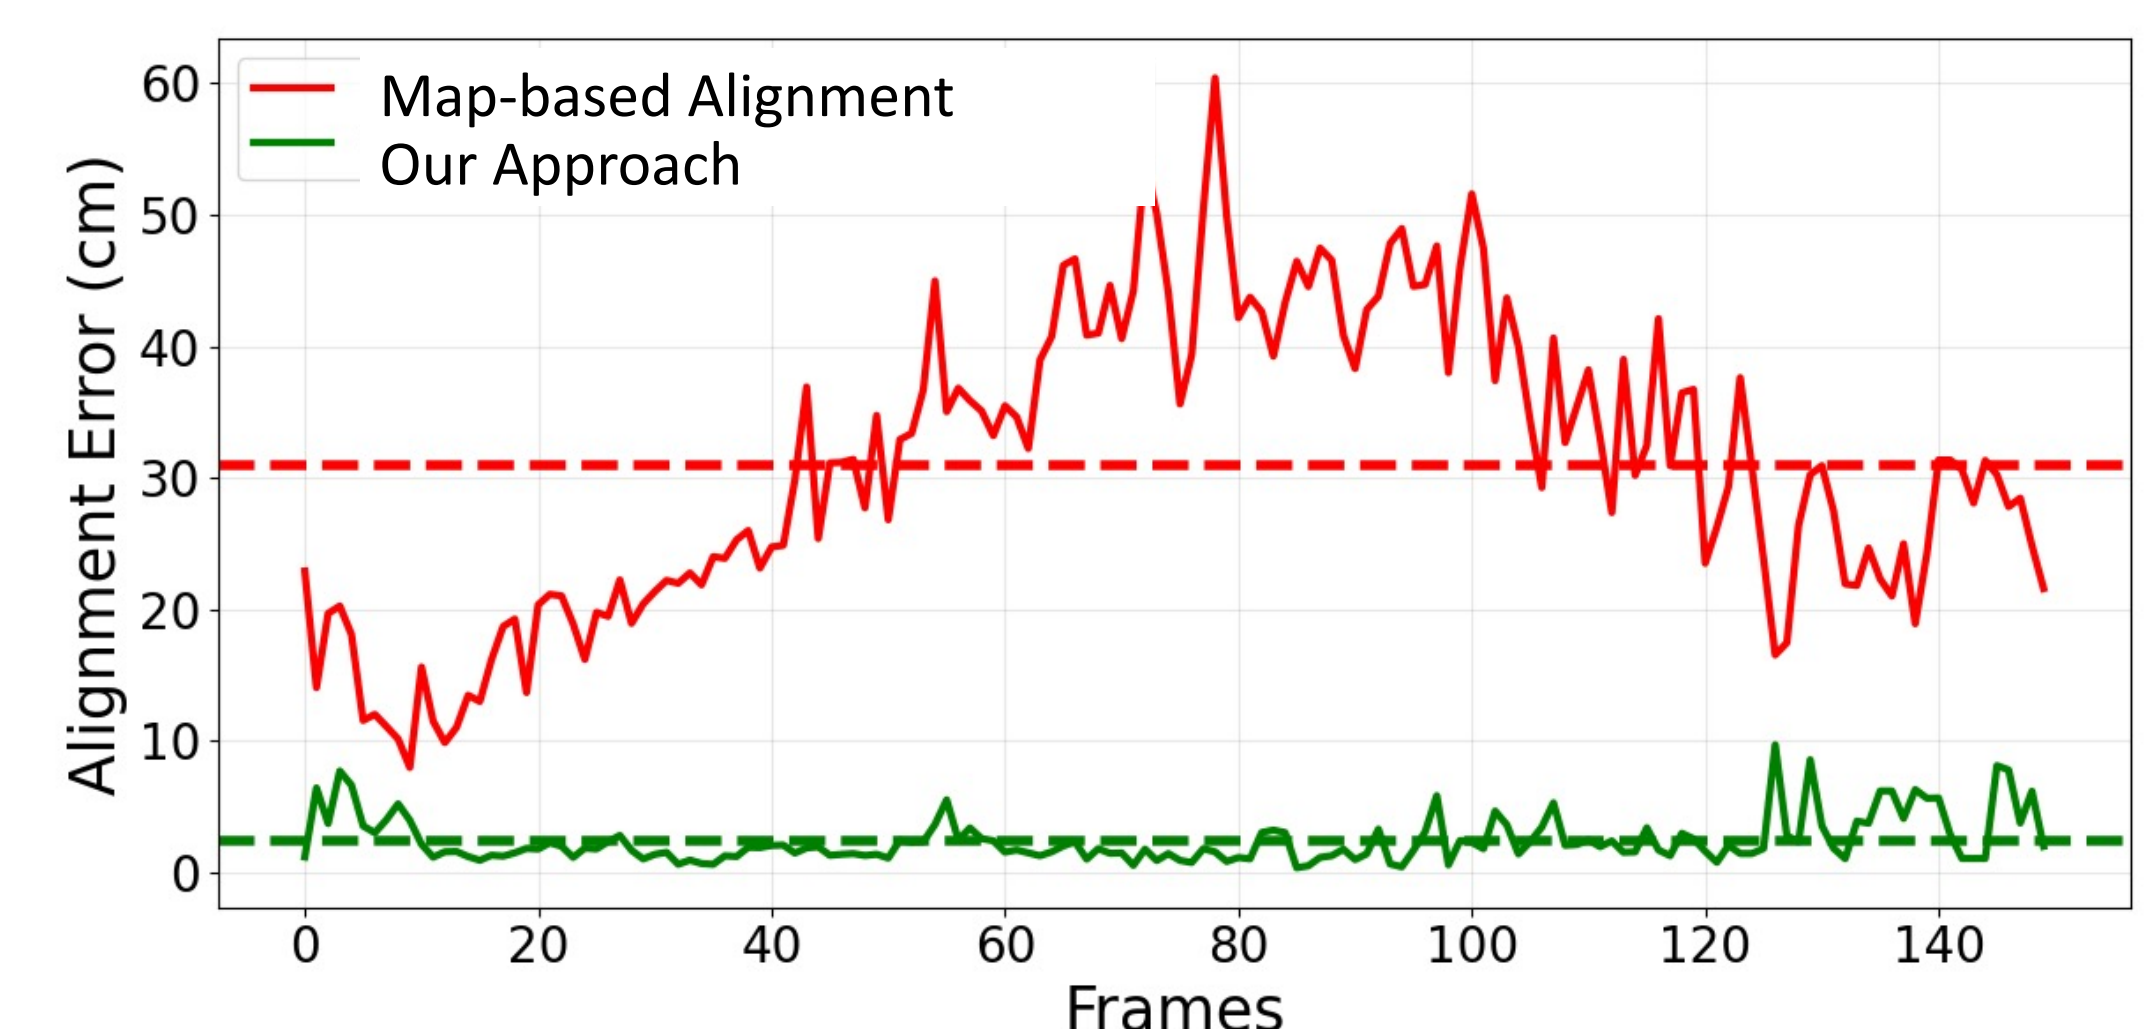


Fig 2: Map-based alignment can be erroneous.




Fig 3: Vehicles directly align themselves to anchors.

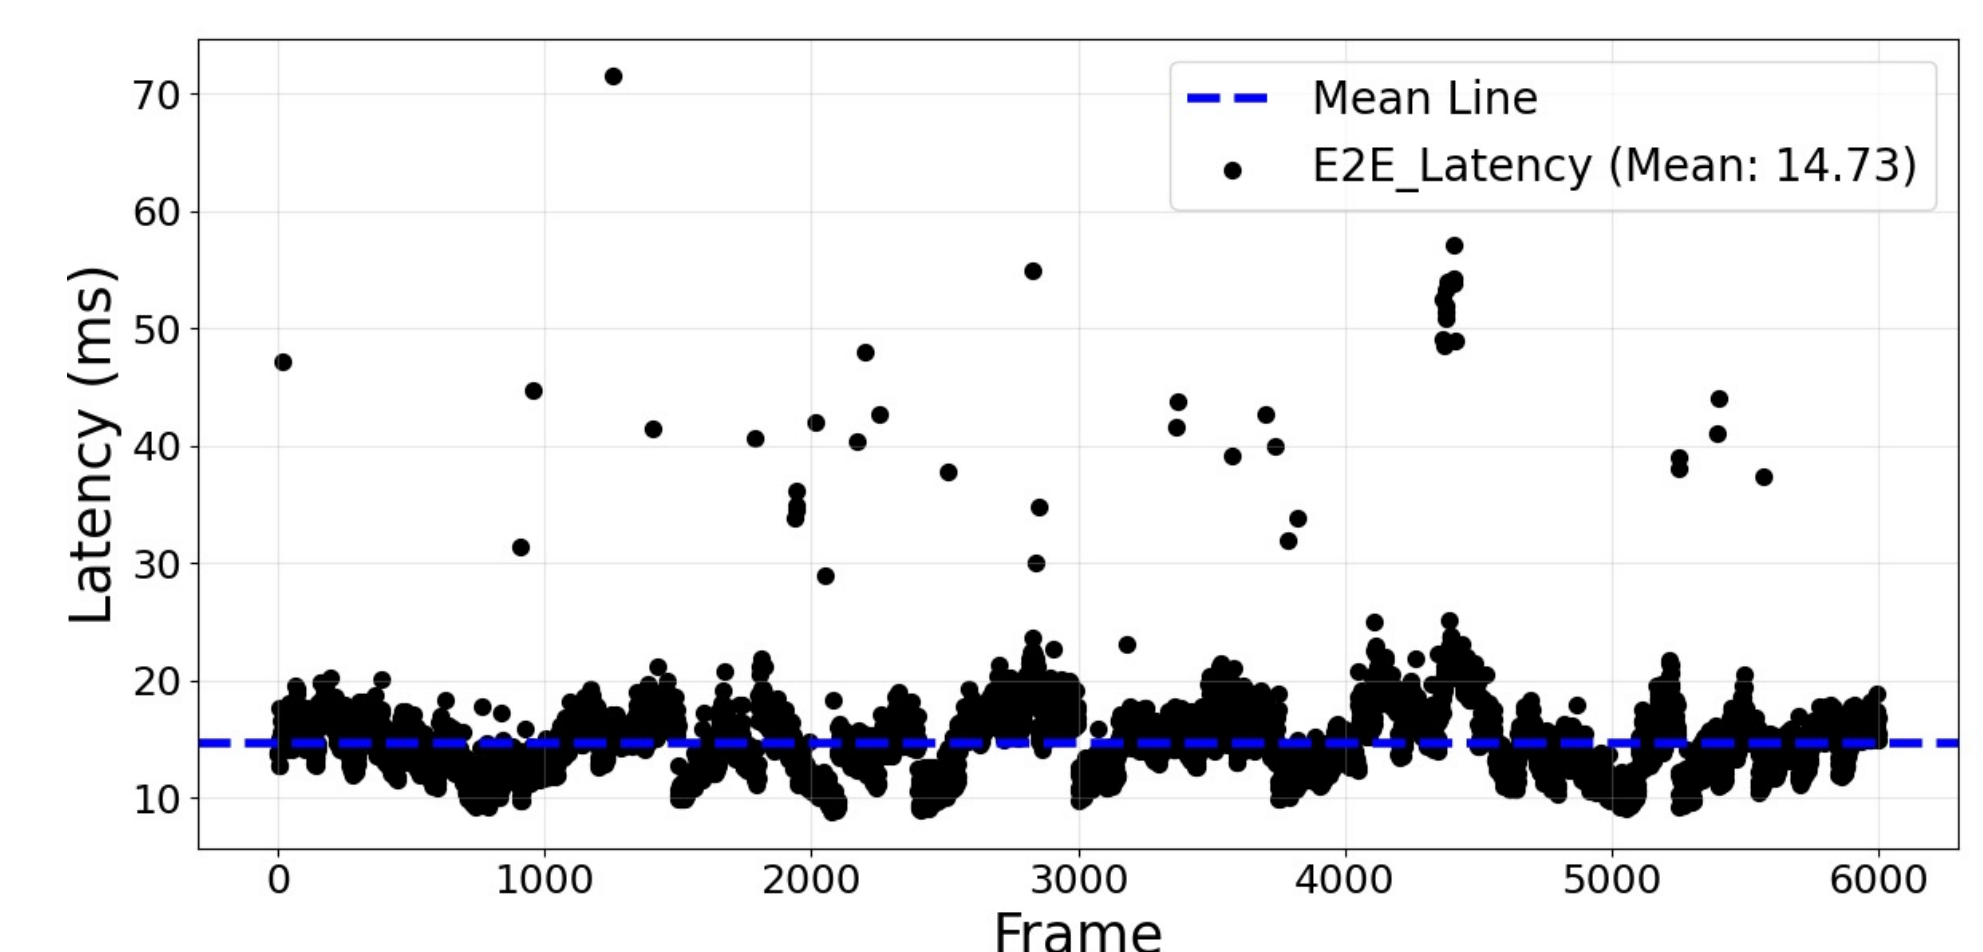


Fig 4: AFS has a mean end-to-end latency of 15 ms.