

Hardware-Accelerated Intelligent Roadside Unit

Hongyi Chen
University of Oxford
hongyi.chen@eng.ox.ac.uk

Noa Zilberman
University of Oxford
noa.zilberman@eng.ox.ac.uk

Introduction A Roadside Unit (RSU) is a multi-functional edge device in intelligent transportation systems (ITS), positioned alongside roadways. Its primary function is to support vehicle-to-everything (V2X) wireless communication, but it can also serve as an intelligent computation, memory, and networking hub [1]. Offloading computation tasks to intelligent RSUs offers many advantages, from eliminating redundant sensors to reducing network traffic [2], but these benefits are diminished by the high processing time on RSUs. To attend to this challenge, we propose a hardware accelerated intelligent RSU, running computation tasks on an FPGA-attached RSU. The hardware acceleration is designed for low-latency, real-time V2X sensor fusion at road intersections, enabling time-critical functions in ITS and improving transportation safety and efficiency.

Motivating Use Case A smart intersection, such as depicted in Fig. 1, is equipped with multiple sensors (such as radars and Lidars) with various ranges. A roadside unit collects the data from all sensors and sensor-instrumented vehicles. The collected information needs to be integrated, which is typically done in the cloud with recent proposals to move to the edge. As different sensors identify objects differently (i.e., no unique identifier), one important integration task is to associate with an object the information from different sensors, and to update it. Typically, this information includes bounding boxes and kinetic information with absolute coordinates and timestamps. In this work, both ID association and data fusion is accelerated in hardware to achieve low-latency.

Why Accelerate? The computation tasks resemble multi-object tracking (MOT) problems to some extent and demand low latency. In the SmartEdge project (www.smart-edge.eu) we aim to achieve end-to-end latency of less than 100ms, much of it consumed by source-sensor and destination-controller. Therefore, ideally an RSU should process messages on millisecond scale. The accelerated algorithm improved and modified the Simple Online and Realtime Tracking (SORT) algorithm, featuring Kalman Filter and Hungarian Algorithm [3]. Benchmarking a python-based implementation of the algorithm on Raspberry Pi 4B, similar in configuration to commercial RSU, yields a latency of about 20 ms per frame, which is an order of magnitude higher than the target, and risky for time-sensitive applications such as collision avoidance. Such a large performance gap is unlikely to be eliminated through software optimising software, demonstrating the need for hardware acceleration.

Implementation Our hardware acceleration solution is




Fig. 1. Smart intersection.

based on the AMD Kria KR260, which uses Zynq UltraScale+ MPSoC. Re-developing the algorithm with hardware description languages enables us to benefit from pipelining, and to reduce latency through parallel object processing and a customized architecture. Trajectory-based datasets, including NGSIM open data and Waymo open motion dataset, are used to generate the data for evaluation, which concentrates on minimizing latency while maintaining the same accuracy as the software-based solution.

Summary This work presents an FPGA-based hardware acceleration solution for intelligent RSUs, integrating bounding boxes and kinetic information of detected and tracked objects collected by different sensors. Future work will incorporate more modalities, for instance, point clouds, to the sensor fusion, and leverage the acceleration outcomes to support time-sensitive applications.

Acknowledgement This work is funded by the EU SmartEdge project (Grant Agreement 101092908, UKRI 10056403).

REFERENCES

- [1] P. Arthurs, L. Gillam, P. Krause, N. Wang, K. Halder, and A. Mouzakitis, "A Taxonomy and Survey of Edge Cloud Computing for Intelligent Transportation Systems and Connected Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6206–6221, Jul. 2022, doi: 10.1109/TITS.2021.3084396.
- [2] J. Zhang, H. Guo, J. Liu, and Y. Zhang, "Task Offloading in Vehicular Edge Computing Networks: A Load-Balancing Solution," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 2, pp. 2092–2104, Feb. 2020, doi: 10.1109/TVT.2019.2959410.
- [3] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple Online and Realtime Tracking," in 2016 IEEE International Conference on Image Processing (ICIP), Sep. 2016, pp. 3464–3468. doi: 10.1109/ICIP.2016.7533003.