報告番号	※甲　　第　　号

主　論　文　の　要　旨

論文題目　GPU-Accelerated Point Cloud Processing for Autonomous Vehicles
（自動運転車向け点群データ処理の GPU 加速）

氏　　名　NGUYEN Anh Viet

論　文　内　容　の　要　旨

Autonomous driving vehicles, or shortly autonomous vehicles can move to their target destination safely with little or no human control. They are expected to reduce human errors, thus increasing driving safety, reducing the cost of transportation, and providing driving experience to people even who cannot drive. To that end, each autonomous vehicle is installed with various kinds of sensors. The sensors collect information from the driving environment in the form of digital data, then pass them to various computing modules to analyze. Decisions on how to control the vehicle are made based on the results of the analyses.

The point cloud is one of the many data types that can be used in an autonomous driving system (ADS). It is generated by laser sensors, such as the Light Detection and Ranging Sensor (LiDARs). With the advance of the technologies, the point clouds generated by LiDARs are becoming more and more detailed, resulting in the increase in their size. Processing such large point clouds is compute-intensive. This poses a challenge for the ADS, which often has strict time requirements.

In this dissertation, we investigate the use of Graphics Processing Units (GPUs) to accelerate point clouds processing in the ADS. We implement a set of fundamental GPU data structures for point cloud data and some GPU-Accelerated primitive algorithms and use them to accelerate compute-intensive applications related to point cloud processings. The three main compute-intensive applications described in this dissertation include: the GPU-Accelerated 3D Normal Distributions Transform, the GPU-Accelerated Euclidean Clustering, and the GPU-Accelerated Indexed Nested Loop Join, which

is one of the heaviest query operations in a Relational Database Management System (RDBMS) for managing point cloud data. The performance of our methods is evaluated via experiments with synthetic and real-world datasets. Based on experimental results, we address the potential and the challenge in developing a GPU-Accelerated sub-system for point cloud processing in the ADS. We also explain some of our future directions to extend our current work toward a fully operational GPU-Accelerated system that can process and manage point cloud data efficiently in the ADS.