Innovative Technology Workshop on 3D LIDAR

Minnesota Department of Transportation

RESEARCH SERVICES & LIBRARY

Office of Transportation System Management

Brian Davis, Principal Investigator
Department of Mechanical Engineering
University of Minnesota

June 2016

Research Project
Final Report 2016-19

LRRB
LOCAL
ROAD RESEARCH
BOARD

Technical Report Documentation Page

1. Report No. MN/RC 2016-19	2.	3. Recipients Accession No.
4. Title and Subtitle Innovative Technology Workshop on 3D LIDAR		5. Report Date June 2016 6.
7. Author(s) Brian Davis		8. Performing Organization Report No.
9. Performing Organization Name and Address Department of Mechanical Engineering University of Minnesota 111 Church Street SE Minneapolis, MN 55455		10. Project/Task/Work Unit No. 2015067 11. Contract (C) or Grant (G) No. (c) 99008 (wo) 197
12. Sponsoring Organization Name and Address Local Road Research Board Minnesota Department of Transportation Research Services & Library 395 John Ireland Boulevard, MS 330 St. Paul, Minnesota 55155-1899		13. Type of Report and Period Covered Final Report 14. Sponsoring Agency Code
15. Supplementary Notes http://www.lrrb.org/PDF/201619.pdf		

16. Abstract (Limit: 250 words)

The use of LIDAR is becoming more common among state, county, and local agencies. It presents a means for collecting a great deal of information about the geometry of a road, its surrounding area, and depending on the sensors used, real-time 3D information about vehicle, cyclist, and pedestrian movements.

The main focus of this project was to develop and conduct two workshops in Minnesota for public DOT and GIS professionals to provide information on the state of the art in mobile LIDAR scanning. Topics included the basics of LIDAR operation, an overview of currently available hardware, as well as current and future applications of the technology. Additionally, the workshops featured a live demonstration of a Velodyne HDL-64E 3D LIDAR scanner.

A sample application was developed to both demonstrate and better understand the capabilities of a real-time 3D LIDAR scanner. This work focused on developing a system capable of automatically collecting vehicle trajectories through intersections using 3D LIDAR data. This application showed that LIDAR might be a suitable tool for collecting traffic data and provided valuable information about the strengths and limitations of such a system.

This project was designed to provide transportation and GIS professionals with accurate, current, and applicable information about LIDAR systems. To accomplish this, existing LIDAR knowledge was combined with market survey research as well as with new information gathered through the process of creating a sample application. This knowledge was aggregated and used to create a workshop that was informative and well received by participants.

17. Document Analysis/Descriptors laser radar, technology transfer, workshops, vehicle trajectories, intersections, traffic counts, turning traffic		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Alexandria, Virginia 22312	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 46	22. Price

Innovative Technology Workshop on 3D LIDAR

Final Report

Prepared by:

Brian Davis
Department of Mechanical Engineering
University of Minnesota

June 2016

Published by:

Minnesota Department of Transportation
Research Services & Library
395 John Ireland Boulevard, MS 330
St. Paul, Minnesota 55155-1899

This report represents the results of research conducted by the author and does not necessarily represent the views or policies of the Minnesota Department of Transportation, the Local Road Research Board, or the University of Minnesota. This report does not contain a standard or specified technique.

The author, the Minnesota Department of Transportation, the Local Road Research Board, and/or the University of Minnesota do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to this report.

ACKNOWLEDGMENTS

This project was funded by the Local Road Research Board. The author would also like to acknowledge the following people and groups for their assistance with this project:

The members of the Technical Advisory Panel (TAP) for their feedback throughout the project and in particular, Kaye Bieniek of Olmsted County who led the TAP as project Technical Liaison.

Dan Sullivan of the Minnesota Department of Transportation, for serving as Project Coordinator and in particular, for assisting by booking venues for the workshops.

Arvind Menon of the University of Minnesota for serving as co-investigator and in particular for the work he performed during the development of the sample application.

Professor Max Donath of the University of Minnesota for his assistance and guidance throughout the project.

Jacob Achtemeier of the University of Minnesota for his assistance in conducting the LIDAR market survey.

The City of Saint Peter, Minnesota for its assistance with the data collection effort.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

Collecting mobile LIDAR data is becoming more common among state, county, and local agencies. It presents a means for collecting a great deal of information about the geometry of a road and its surrounding area at a lower cost than traditional surveying techniques. These methods yield a point cloud, or 3D geometric representation, of the scanned area. In addition to capturing static information about road geometry, real-time 3D scanners are also capable of generating many complete point clouds per second, providing information about not only the road itself but also about the behaviors of the vehicles, cyclists, and pedestrians using it. Depending on the accuracy of the sensors used, the resulting point cloud data can be accurate to 1 cm or better. Such systems enable a number of applications including surveying, construction, driver assistance systems, facilities inspection, inventory detection, asset management, and more.

The main focus of this project was to develop and conduct two workshops in Minnesota for public DOT and GIS professionals to provide information on the state of the art in mobile LIDAR scanning. Topics included the basics of LIDAR operation, an overview of currently available hardware, as well as current and future applications of the technology. Additionally, the workshops featured a live demonstration of a Velodyne HDL-64E 3D LIDAR scanner.

A sample application was developed to both demonstrate and better understand the capabilities of a real-time 3D LIDAR scanner. The application selected for development was the design of a system capable of automatically collecting vehicle trajectories through intersections using 3D LIDAR data. The sample application development yielded valuable information about the strengths and limitations of such a system and provided insight into the identification of future directions to take the application. In short, this application showed that LIDAR might be a suitable tool for collecting traffic data.

This project was designed to provide state, county, and local transportation and GIS professionals with accurate, current, and applicable information about LIDAR systems. To accomplish this, existing LIDAR knowledge was combined with market survey research as well as with new information gathered through the process of creating a sample application. This knowledge was aggregated and used to create a workshop that was informative and well received by participants.

CHAPTER 1: INTRODUCTION

1.1 Project Overview

Collecting mobile LIDAR data is becoming more common among state, county, and local agencies. It presents a means for collecting a great deal of information about the geometry of a road and its surrounding area at a lower cost than traditional surveying techniques. These methods yield a point cloud, or 3D geometric representation, of the scanned area. In addition to capturing static information about road geometry, real-time 3D scanners are also capable of generating many complete point clouds per second, providing information about not only the road itself but also about the behaviors of the vehicles, cyclists, and pedestrians using it. Depending on the accuracy of the sensors used, the resulting point cloud data can be accurate to 1 cm or better. Such systems enable a number of applications including surveying, construction, driver assistance systems, facilities inspection, inventory detection, asset management, and more.

The main focus of this project was to develop and conduct two workshops in Minnesota for public DOT and GIS professionals to provide information on the state of the art in mobile LIDAR scanning. Topics included the basics of LIDAR operation, an overview of currently available hardware, as well as current and future applications of the technology. Additionally, the workshops featured a live demonstration of a Velodyne HDL-64E 3D LIDAR scanner.

In addition to the workshop material development, a sample application was developed using the Velodyne HDL-64E. The goal of this effort was to provide valuable hands-on experience with the sensor and its software suite in a real-world application. The application selected was to automatically generate vehicle turn counts for vehicles traveling through an intersection. This particular application was chosen because the results represent realistic data needed by traffic professionals. Furthermore, the application and supporting software development would facilitate an evaluation of the hardware and available software tools which would help inform the material for the workshop.

1.2 Workshop Format and Agenda

Two equivalent workshop sessions were held. The first was hosted by Dakota County at their Northern Service Center in West St. Paul. The second was hosted by St. Louis County in their Public Works and Transportation Complex in Duluth. These locations were chosen to maximize accessibility for potential participants throughout the state. Care was also taken to locate venues that could accommodate both a classroom component as well as live LIDAR demo.

The workshop agenda was split into three main sections. The first section was an overview of LIDAR to help familiarize participants with the technology by first explaining the operating principles, defining the vocabulary used when discussing LIDAR, and discussing the significant specifications of LIDAR scanners as they are used to differentiate sensors. An overview of LIDAR sensors from different application areas (aerial, survey, mobile) was presented and discussed to show the range of hardware currently available and the applications they enable. Additionally, an overview of other commonly used on-vehicle sensors was presented covering, a brief introduction to GNSS/GPS, RADAR, and vision based systems. The goal of this information was to provide context and help describe the landscape of technologies with which LIDAR is frequently integrated.

The second section of the workshop was a live demonstration of the Velodyne HDL-64E. The sensor was mounted so it could capture point cloud data of its surroundings and send that information to a nearby computer that showed a real-time visualization of the collected data. This component of the workshop was designed to allow for participants to get a better qualitative sense of both the LIDAR's physical characteristics and how it operates including the relative point densities for objects at different distances, how the sensor responds to objects with differing retroreflectivities, and how object occlusion affects the returns.

The last section of the workshop consisted of an in-depth look at two sample applications of LIDAR technology. The first was a previously conducted project that investigated using a side-facing planar LIDAR scanner mounted on a vehicle to determine the presence and position of roadside assets such as curbs and guard rails. The second example discussed was the sample application developed as a part of this project that used a LIDAR to determine vehicle turn counts through an intersection.

1.3 Report Organization

This report documents the work performed as a part of this project including the market survey research and the sample application development. Chapter 1 provides an overview of the project and information about the workshops themselves.

Chapter 2 is a summary of the first section of the workshop which focused on the current state of LIDAR. This chapter includes the material presented about the LIDAR's operating principles, major specification definitions as well as an overview of currently available hardware. This report will not include the information presented on the complimentary technologies (GNSS/GPS, RADAR, vision systems) as information about these systems was not the focus of this project.

Chapter 3 discusses the vehicle counting application developed as a part of the project and presented at the workshops. It includes the system design, algorithms used to detect and track the vehicles, data collection, and the results of the experiment. It also summarizes the strengths and limitations of the developed system (both hardware and software) as it was implemented. This report will not include a discussion of the other application discussed at the workshops (high accuracy mapping of roadside assets) as information about this work is available in the final report for that project [1].

Chapter 4 summarizes the significant findings of the research as well as feedback collected at the workshops.

CHAPTER 2: CURRENT STATE OF LIDAR

2.1 Introduction to LIDAR

LIDAR (LIght Detection And Ranging) is a technology that uses laser pulses to determine the geometry of the space around a sensor. It has a number of applications in many different industries including mapping, surveying, asset management, and advanced driver assistive systems (ADAS) among others.

LIDAR determines the positions of the geometry around it by emitting laser pulses that reflect off of surrounding objects and return to the sensor. It times how long it takes for the pulse to travel to the object, reflect, and return to the sensor and uses that time to calculate how far away the object is located. Scanners construct an entire point cloud by rotating the emitter (or an array of emitters) so that pulses get sent in many directions.

Different LIDAR scanners vary in a number of ways ranging from their accuracy, the area or volume they can scan, how much information is collected and based on this, how they must be mounted to be useful for a particular application. Another consideration is what software tools are available with which to interface and collect data from the hardware so that it is presented or processed in a useful way.

Accuracy is an overall measure of how close the generated point cloud is to the actual positions of the surfaces within a scanned scene. Distance accuracy generally refers to how accurately a sensor can determine the distance from the sensor to a particular point. Additionally, point cloud accuracy (also called lateral or angular accuracy) is a measure of how close an individual point is within the point cloud when compared to the actual position that the pulse hit. This accuracy is most important when examining the error associated with the absolute position of the object's surface with respect to the sensor's location.

Sensor range is an important consideration when determining a sensor's suitability for a scanning task. Obviously the geometry to be scanned must be within the sensors maximum range. Additionally, the maximum range is affected by the reflectivity of the target. Objects that are very retroreflective (i.e. they return the pulses back to the sensor with minimal intensity loss), are easier to detect at further distances since more of the pulse's energy is returned to the sensor. Objects that are reflective, but not necessarily retroreflective may be problematic for the sensor. For example, a mirror is very reflective but depending on its orientation, it doesn't necessarily return the pulse back to the sensor. The pulse may reflect off of the mirror and never return to the sensor.

Point density is a combination of a sensor's vertical and horizontal angular resolution and frequency. These factors together help describe how many points make up a particular object's surface within the point cloud. Because the pulses diverge as they move away from the sensor, objects that are further from the sensor are represented with fewer points than those that are closer. This means that in addition to considering the maximum range of the sensor, the point density must be high enough that the objects to be detected are also represented with enough information to be useful.

Field of view is a measure of the area or volume where the sensor can detect objects. It is described by the vertical and horizontal field of view. LIDAR range from 1 dimension (a single pulse range finder) to 2D units that scan in a single plane to 3D units that can scan an entire volume around the unit. Consider the Velodyne HDL-64E [2] which has a vertical field of view of roughly 28° spanning from 2° above level to 26° below level. Here, level refers to the plane that intersects the emitter and is parallel to the base (regardless of the sensor's orientation with respect to the ground). Because the housing itself spins about its base, this unit has a 360° horizontal field of view. This is illustrated in Figure 2.1.

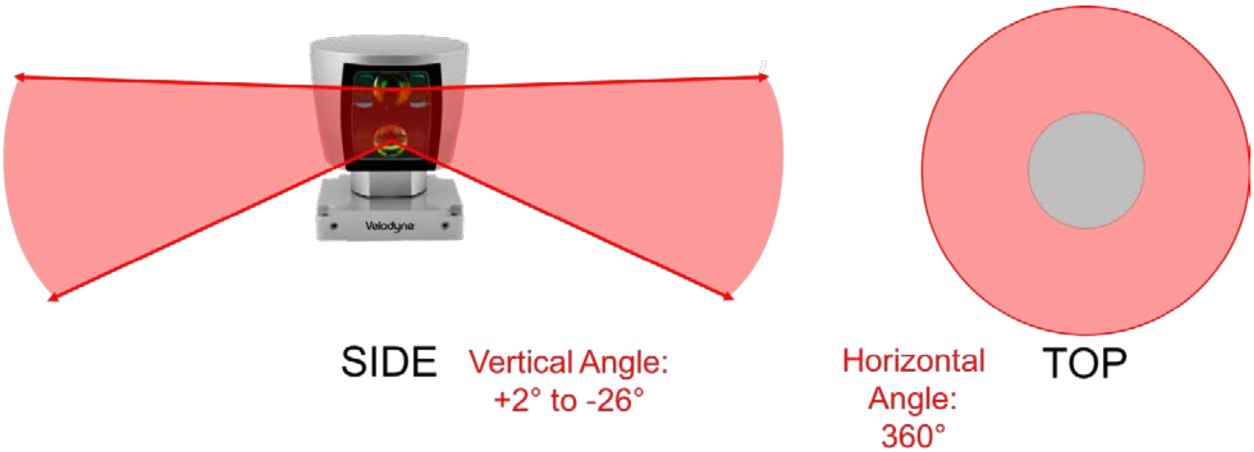

Figure 2.1 Illustration of Velodyne HDL-64E field of view.

Obviously the sensor can't detect objects outside of its field of view. This manifests itself as a dead zone or shadow beneath the sensor. This can be problematic as the sensor is mounted higher off the ground because the higher the sensor is, the larger the cast shadow's radius. For example, when this sensor is mounted on top of a vehicle, it is approximately 2 meters above the ground. This results in a shadow with a radius of approximately 4 meters.

Planar LIDAR are sensors that instead of scanning a volume, only scan in a single plane. They are capable of providing only 2 dimensions of information about the plane they are scanning. For example, the Sick LMS 511 [3] is a scanner that has a horizontal field of view of 190°. Because it scans in a plane, the effective vertical field of view is 0°. This is illustrated in Figure 2.2.

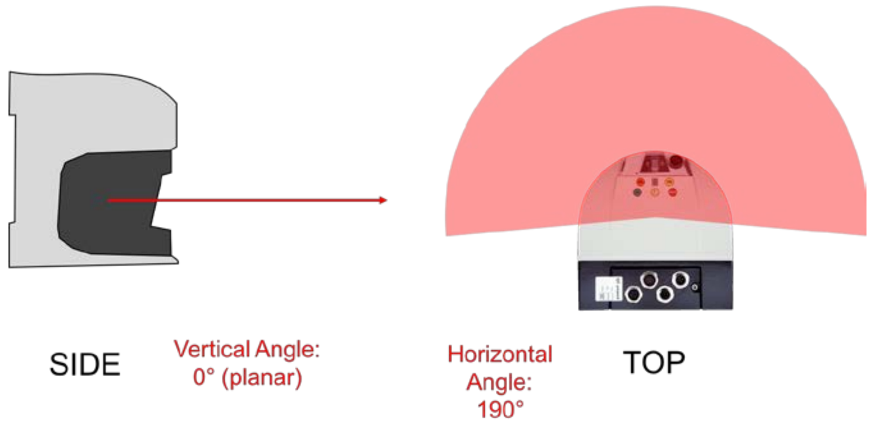

Figure 2.2 Illustration of Sick LMS 511 field of view.

Software is another significant consideration when evaluating LIDAR. Software is required to interface with the sensor to parse the raw serial data output and send configuration commands. Generally, data visualization will be important to understand what the sensor is observing. Data must be captured or recorded and played back to analyze previously collected information. Then, depending on the application, either real-time processing or post-processing will need to occur to provide useful information based on the point cloud data collected by the sensor.

Software tools range from custom software to end-to-end software packages provided by the vendor. Some open source tools exist for visualizing, recording, playing back data as well as some processing tasks like object identification and tracking. However, implementing these solutions still requires a great deal of development to leverage them in a useful piece of software.

2.2 Overview of Available Hardware and Current Applications

LIDAR technology is used for numerous applications across a diverse collection of industries. This overview provides a brief summary that highlights some currently available LIDAR hardware along with the applications they are most commonly used for. For the purposes of this overview, hardware will be presented as belonging to one of three categories: aerial LIDAR, stationary surveying LIDAR, and mobile LIDAR. These distinctions are made only to help frame the hardware within their applications. In reality, many of the technologies described bridge the gap between these categories.

In order to provide accurate and complete information, this overview is comprised of information about real, currently available LIDAR systems. The authors would like to re-iterate that this information does not constitute an endorsement of any brand or product on behalf of the authors, the University of Minnesota, the Local Road Research Board, or the Minnesota Department of Transportation. Unless otherwise noted, all specifications and features are provided as presented by the manufacturer in publicly available press material, data sheets, and product manuals.

2.2.1 Aerial LIDAR

Aerial LIDAR is the use of airborne LIDAR scanners mounted on aircraft to characterize the surface of the earth. Typically, the most common use is the generation of digital terrain models for elevation. However, it can also be used to collect data about man-made objects such as buildings or roads.

The systems used for aerial LIDAR generally combine LIDAR scanning technology, which provides information about the geometry of the ground, with positioning instruments such as Global Navigation Satellite Systems/Global Positioning System (GNSS/GPS), inertial measurement units (IMUs), or other positioning technology to determine the position and path of the airplane. When combined, they are capable of generating an accurate representation of the earth's surface.

The Trimble AX60 [4] integrates a Riegl LMS-Q780 LIDAR with Trimble positioning equipment to create a complete aerial LIDAR product. It is capable of producing a point cloud with vertical accuracy of better than 15 cm, horizontal accuracies better than 20 cm and the system scans at a rate of 200 Hz. This package also includes Trimble's software used to interface with and operate the hardware as well as analyze and transform the data when back on the ground. Figure 2.3 shows an image of the hardware.

Figure 2.3 Trimble AX60 aerial LIDAR system.

Another example is the Optech Orion [5] aerial LIDAR. This hardware is similar but here the vendor specifications report the accuracies as a range noting that the altitude at which the airplane is flown affects the accuracy. Lower altitudes allow for more accurate, denser point clouds but at the cost of collecting less data per pass. The vertical accuracy is 3 to 15 cm and the horizontal accuracy is 2 to 50 cm. Figure 2.4 shows the system as installed in an airplane before flying for a project in southeastern Minnesota.

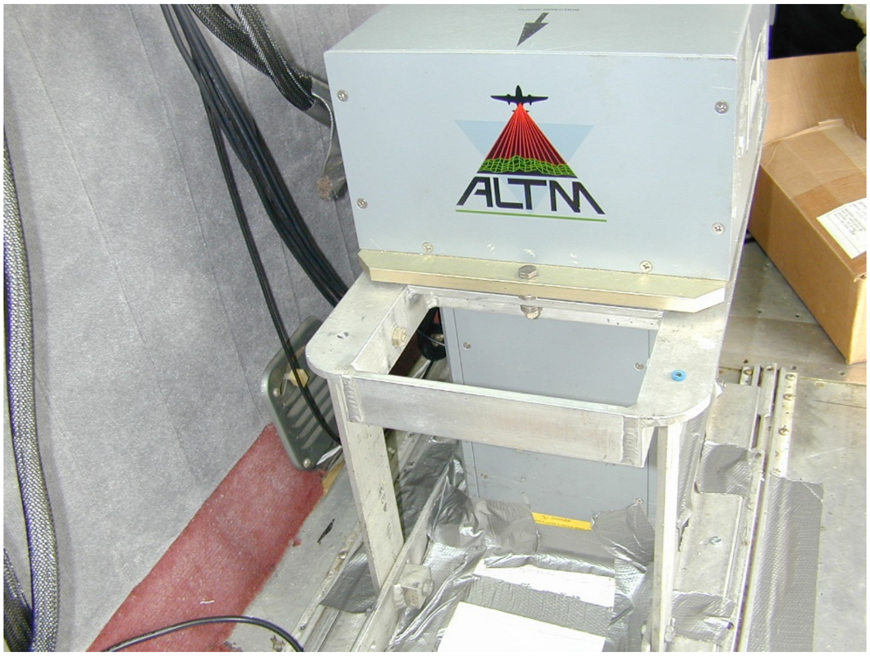

Figure 2.4 Optech Orion aerial LIDAR system mounted in airplane.

2.2.2 Stationary LIDAR

Stationary LIDAR is a segment that focuses on creating high accuracy, static point clouds for an area. Generally, these systems are mounted on a base and slowly move to scan the area around it creating a single, information-rich point cloud. These types of systems have many applications including surveying and construction, forestry and agriculture, among others.

The Riegl VZ-2000 [6] is one such LIDAR system. It has a 100° vertical field of view and is designed to rotate such that it has an effective 360° horizontal field of view. It's capable of scanning at a maximum range of 2000 meters with an accuracy of 8 mm. The system is shown in Figure 2.5.

Figure 2.5 Riegl VZ-2000.

In addition to the tripod mounting configuration, the manufacturer also provides a mount so the LIDAR may be attached to a car. This bridges the gap between stationary LIDAR and mobile LIDAR scanning discussed in the next section. The system mounted on a vehicle is shown in Figure 2.6.

Figure 2.6 Riegl VZ-2000 mounted on a vehicle.

Another example of stationary LIDAR technology is the Optech ILRIS [7]. It is a unit capable of scanning at a maximum range of 2000 to 3000 meters. The measurements have a 4 mm range accuracy and an 8 mm angular accuracy. The scanner only has a 40° vertical by 40° horizontal field of view. However, this unit also has a motorized base that is capable of both rotating side-to-side and tilting up and down for a much greater effective field of view. The unit is shown in Figure 2.7.

Figure 2.7 Optech ILRIS.

An example scan of a boat lift collected by the hardware is shown in Figure 2.8.

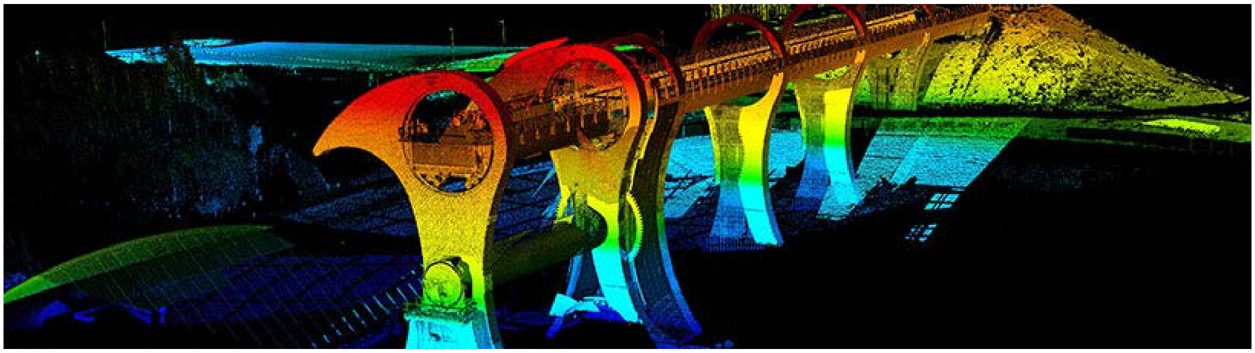

Figure 2.8 Optech ILRIS scan of boat lift.

2.2.3 Mobile LIDAR

Mobile LIDAR is the use of LIDAR technology on terrestrial vehicles. It's used to capture information about a vehicle's surroundings including the road geometry, roadside obstacles, other vehicles, and more. This data can be used in many different ways from creating maps to real-time use in autonomous vehicle systems. Mobile LIDAR captures information about a corridor that is more detailed than aerial LIDAR but is faster to collect than setting up multiple scans with a stationary LIDAR product.

The first example is a 1 dimensional LIDAR unit from PulsedLight called the LIDAR-LITE v2 [8]. It is capable of scanning in only a single direction, finding the distance to the nearest object at which it's pointed. In this capacity, it can be thought of as a type of range finder. The unit costs roughly $115 and is capable of a 40 m range with a 2.5 cm accuracy. The unit is shown in Figure 2.9.

Figure 2.9 PulsedLight LIDAR-LITE v2.

Units like this have applications as security system components, industrial fluid/solid level measurements, and for collision avoidance. One such use is as a collision avoidance warning system for bicycles [9]. The sensor is mounted on the back of a bicycle facing backwards. The distance measurement provided is then used to determine when a vehicle is approaching the bicycle at an unsafe rate so an alarm can be sent to the cyclist. This configuration is shown in Figure 2.10.

Figure 2.10 PulsedLight LIDAR-LITE v2 mounted on a bicycle.

The Sick LMS 511 [2] is an example of a planar LIDAR. It scans in a single plane with a horizontal field of view of 190°. It's capable of 1/6° angular resolution at a 25 Hz scan frequency. The overall accuracy of the data it produces is 12 mm. These types of scanners have applications in transportation, industrial object flow, and security. The sensor is shown in Figure 2.11.

Figure 2.11 Sick LMS 511.

The IBEO ScaLa [10] is an example of a “near”-planar LIDAR. This means that although there is a limited vertical field of view, it is not designed to create fully 3D point clouds. This unit has a 145° horizontal field of view and a 4 plane, 3.2° vertical field of view. This sensor is designed for use in advanced driver assist systems for sensing objects around a vehicle. The vertical field of view allows for robustness to a vehicle’s changes in pitch as it navigates a roadway. The unit has a 150 m range and a 10 cm distance resolution. Additionally, the hardware also features embedded feature/target tracking. This means that in addition to outputting the raw point cloud, on-sensor processing is also performed to provide the position of individual objects. This sensor is shown in Figure 2.12.

Figure 2.12 IBEO ScaLa.

The Riegl VQ-450 [11] is a high-end planar LIDAR capable of producing highly accurate and dense information about the area it scans. It has a horizontal field of view of 360° and an accuracy of 8 mm. It is capable of generating 150,000 to 550,000 points per second at up to 200 scans per second. The angular step width (angle between sequential pulses) is variable from 0.48° to as small as 0.001°. These units are relatively expensive and generally are used for high accuracy applications such as road or corridor surveying. An image of the sensor is shown in Figure 2.13.

Figure 2.13 Riegl VQ-450.

Using LIDAR for surveying is a frequent use case for higher end scanning systems. Generally, this application utilizes scanning platforms that integrate LIDAR with positioning systems such as GNSS/GPS, IMUs, and distance measurement indicators (DMIs). The example discussed below uses planar LIDAR scanners that are not by themselves, capable of creating a 3D point cloud. Rather, they rely on the vehicle’s motion so that their field of view sweeps past or travels through the corridor to be digitized.

One such example is the RIEGL VMX-450 [12]. It utilizes two VQ-450 LIDAR scanners and integrates them with the positioning technology required to create geo-referenced point clouds. The system is capable of 2 to 5 cm absolute position accuracy that represents the offset between an object's real position and the position data collected by the system. However, the relative position accuracy can be as good as 1 cm or better. This is a measure of the offset between the distance between two objects' real positions and the same distance as measured in the point cloud. The system is shown as mounted on a vehicle in Figure 2.14.

Figure 2.14 Riegl VMX-450 mounted on a vehicle.

Real-time 3D LIDAR is a type of LIDAR that is focused on providing a full 3D point cloud that is updated many times per second. This is in contrast to other systems discussed that generate a full point cloud by slowly revolving or that require a vehicle to move to collect corridor information. This type of sensor can be used either in a mobile or stationary configuration and can enable applications such as mapping, surveying, and autonomous vehicle guidance.

One such LIDAR is the Velodyne HDL-64E [2] which is capable of scanning with 64 channels at once. This means that over its 26.8° vertical field of view, it emits 64 pulses at once. The housing of the unit spins about its base which allows for a 360° horizontal field of view. This enables a complete point cloud to be created each time the LIDAR completes one revolution. The resulting data is accurate to the 2 cm level or better and the system generates 2.2 million points per second. The sensor is shown in Figure 2.15.

Figure 2.15 Velodyne HDL-64E.

Two samples of collected point cloud data is shown in Figure 2.16. Note that in these visualizations the location of the LIDAR is above the center of the black circle which corresponds to the dead zone or shadow beneath the sensor.

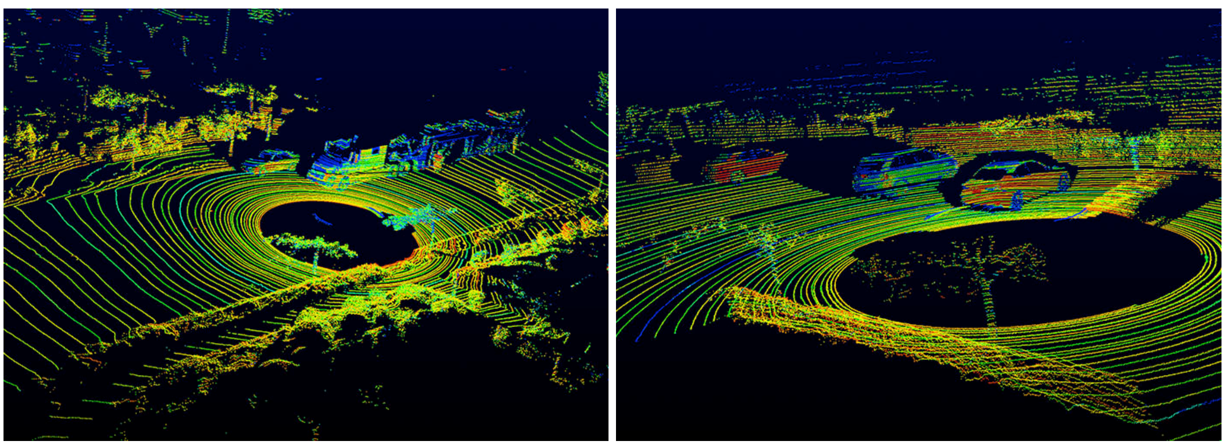

Figure 2.16 Sample point cloud data from Velodyne HDL-64E.

Velodyne also makes two smaller units, the HDL-32E [13] and the Puck [14]. These are units that scan with 32 channels and 16 channels respectively. They have similar accuracies and ranges, the major difference being the density of the point clouds collected by the sensors.

CHAPTER 3: VEHICLE COUNTS THROUGH INTERSECTIONS

3.1 Introduction

A sample application was developed to both demonstrate and better understand the capabilities of a real-time 3D scanning LIDAR. The application selected for development was the design of a system capable of automatically collecting vehicle trajectories for vehicles passing through an intersection using 3D LIDAR data. This system was designed to operate by mounting a Velodyne HDL-64E LIDAR on the top of a vehicle and then positioning the vehicle near an intersection. Then the LIDAR's data could be recorded for later post-processing by an algorithm capable of determining trajectories for individual vehicles as they pass through the intersection. This would allow for the software to aggregate this data and report turn counts.

There were two main goals of this effort. The first was to explore new uses for 3D LIDAR in applications that would be useful for traffic professionals. Vehicle turn counts at intersections represent a commonly collected type of data that city and county traffic engineers frequently need to generate. Current methods for collecting this data generally require infrastructure to be deployed such as tube counters or manual hand coding of traffic behavior either on-site or from recorded video, which is effort intensive. The proposed method would allow for minimal human intervention necessary for intersection setup but once the vehicle with the LIDAR is parked at an intersection, almost no effort would be required other than occasionally monitoring the system to ensure normal operation.

The second goal for this work was to examine the hardware and included software to identify their strengths and limitations. It provided an opportunity to gain valuable hands-on experience with the system allowing for the examination of operating specifications, ease of use, software extensibility, and reliability. This would allow for more informed, experiential content for the workshop that would be more valuable for attendees.

3.2 Proposed System

3.2.1 LIDAR

This application used a Velodyne HDL-64E S2 [2] LIDAR. It's a real-time 3D LIDAR scanner configured to emit 64 laser pulses at once spread across its vertical field of view which spans from approximately 2° above level to 25° below level. The entire upper section of the housing rotates about its base to capture a full 360° horizontal field of view. This results in a scanned volume that is mostly level or below the sensor. The unit is shown below in Figure 3.1.

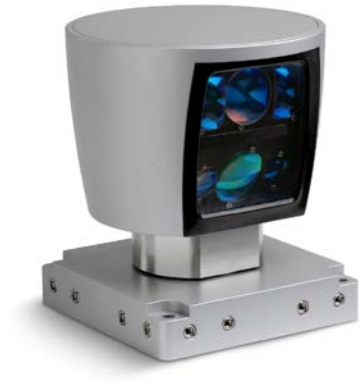

Figure 3.1 Velodyne HDL-64E.

This scanner is designed to spin at a scan frequency between 5 and 15 Hz (revolutions per second), with each revolution generating a complete 3D point cloud of the surrounding area. Depending on the scan frequency, the angular resolution of the data collected varies between 0.09° and 0.35° (slower revolution speeds allowing for denser data). This results in approximately 1.333 million points being collected per second with an overall point cloud accuracy of 2 cm.

3.2.2 Mounting Hardware

Mounting hardware was designed and machined so that the LIDAR could be attached securely to the top of the research vehicle, a Chevrolet Impala. The hardware consisted of two main components. The first was an off-the-shelf roof rack mounting kit that was used to mount two solid aluminum rods such that they spanned the roof over the rear passenger doors. The kit's stock, hollow mounting pipes were replaced with these solid aluminum rods for added strength to help reduce vibrations or bouncing when supporting the weight of the LIDAR.

The second component was a metal mounting plate that attaches the LIDAR to the rods. The plate has a hole pattern that allows for the LIDAR to bolt to the plate and slots on the plate that allow for U-bolts to connect the plate to the rods. When fully assembled and mounted, the LIDAR's base is roughly 4.5 inches above the roof of the car and the top of the LIDAR is positioned roughly 6 feet above the ground. This configuration is shown in Figure 3.2.

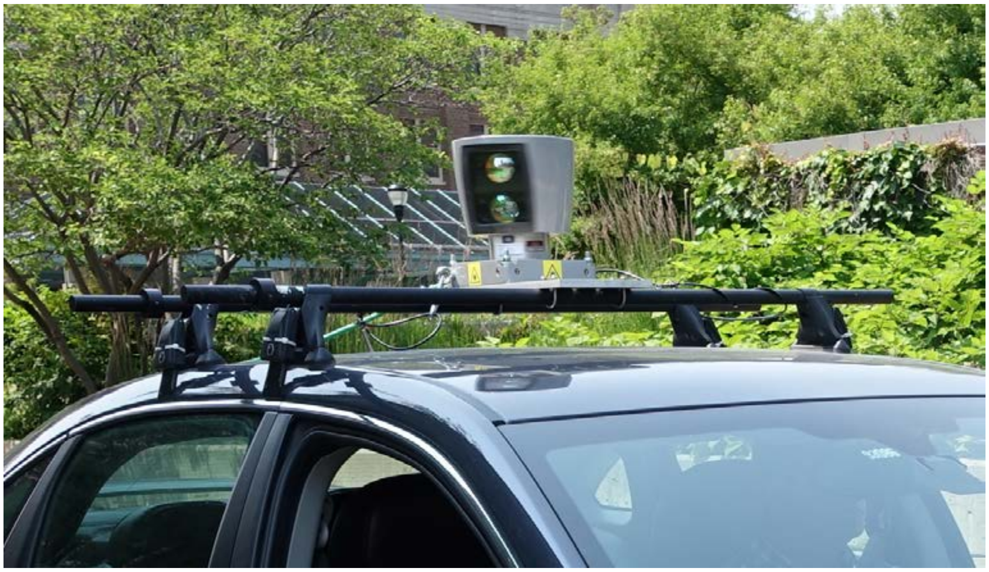

Figure 3.2 Velodyne HDL-64E LIDAR mounted on research vehicle.

This mounting configuration was selected because it allowed for a relatively inexpensive, straight forward mounting solution. It also balanced the height of the sensor above the roof with overall sensor height which was a consideration due to logistical concerns associated with fitting the vehicle in areas with limited vertical clearance.

3.2.3 Other Hardware

In addition to the LIDAR itself and the mounting hardware, a computer was used to run the software to visualize and capture the data. The computer used was a consumer grade laptop with an Intel Core i7 CPU. This computer required a high performance CPU to aid in post-processing tasks, but the system does not require real-time processing. This means that a computer with lower performance hardware would likely still be suitable for this task.

The LIDAR provides an Ethernet interface for data communications and a terminal block to receive power. Additionally, there is also a serial interface for issuing configuration commands to the LIDAR but that connection was unused after a one-time initial setup. The computer and LIDAR were attached to the same network switch which provided a data link between them. The terminal block was wired to an existing power bus in the trunk of the car which supplies 12V DC power. The LIDAR accepts this voltage without need for a separate power supply.

3.2.4 Data Acquisition Software

The software recommended by the manufacturer to interface with the LIDAR is an open source package called VeloView. It is used to interface with the LIDAR and parse the raw data stream it sends to convert it into a point cloud which consists of x,y,z coordinates for each observed point. The software is also capable of providing a 3D visualization of the point cloud which is useful to confirm the system is operating correctly, that the field of view contains the area where data is to be captured, among other uses.

The software is also capable of recording live data to store it for later use. The software captures the raw bit stream being sent from the LIDAR as this is more space-efficient than storing a table of x,y,z points for each frame. Then, as the data is played back the software re-parses the data from the file as if it were being sent from a live LIDAR.

3.3 Data Collection

LIDAR data was collected at two intersections to be used for the development and testing of the post-processing algorithms designed to determine vehicle trajectories and return vehicle turn counts. The two intersections were in southern Minnesota and were selected based on the following criteria. First, the intersections needed to have a safe location where the vehicle could be stationed to collect data and the research team needed to be able to secure permission to conduct the data collection effort. The second criterion was that the intersections would represent an edge case for the system or would otherwise be interesting such that system limitations could be characterized and analyzed.

Data was collected in 10 minute segments. This was done to limit the size of the files and to reduce the risk of losing data due to file corruption or other unforeseen issues. Even by breaking up the data collection in this way, individual files containing 10 minutes of data were roughly 2.6 GB when stored. Seven of these 10 minute files were collected at each intersection over the course of roughly 90 minutes. Both data collection efforts were conducted on a day with clear weather in late September 2016.

The first data collection effort was in Rochester, MN at the intersection of 16th St SW and Mayowood Rd. Here, 16th St SW is a major, undivided 4 lane road that does not have traffic control in either direction. Mayowood Rd is a smaller, undivided 2 lane road that meets 16th St SW at a T intersection. Vehicles approaching on Mayowood Rd have a stop sign. This intersection only has one protected turn lane which is for vehicles making a right turn from 16th St SW to Mayowood Rd. The intersection is located near US Highway 63 on the southern side of Rochester. Its location shown in Figure 3.3.

Map data © 2016 Google

Figure 3.3 Rochester intersection location.

The research vehicle (with the LIDAR mounted on top) was positioned alongside 16th St SW in a private driveway opposite where Mayowood Rd meets 16th St SW. Satellite photography of the intersection is shown in Figure 3.4 which also notes the location where the LIDAR was stationed.

Imagery © 2016 Google, Map data © 2016 Google

Figure 3.4 Rochester intersection satellite photography.

This intersection was selected due to a number of characteristics. First, it was a relatively large intersection. The main road has 4 thru lanes and a protected right turn for vehicles exiting to the side road. It would also provide a better test of a road with higher volume of traffic. Lastly, the geometry of the intersection was such that the LIDAR could be stationed directly across from where the side road meets the main road in a driveway which would allow for a safe, out of the way deployment.

Figure 3.5 shows the intersection from the view of the research vehicle.

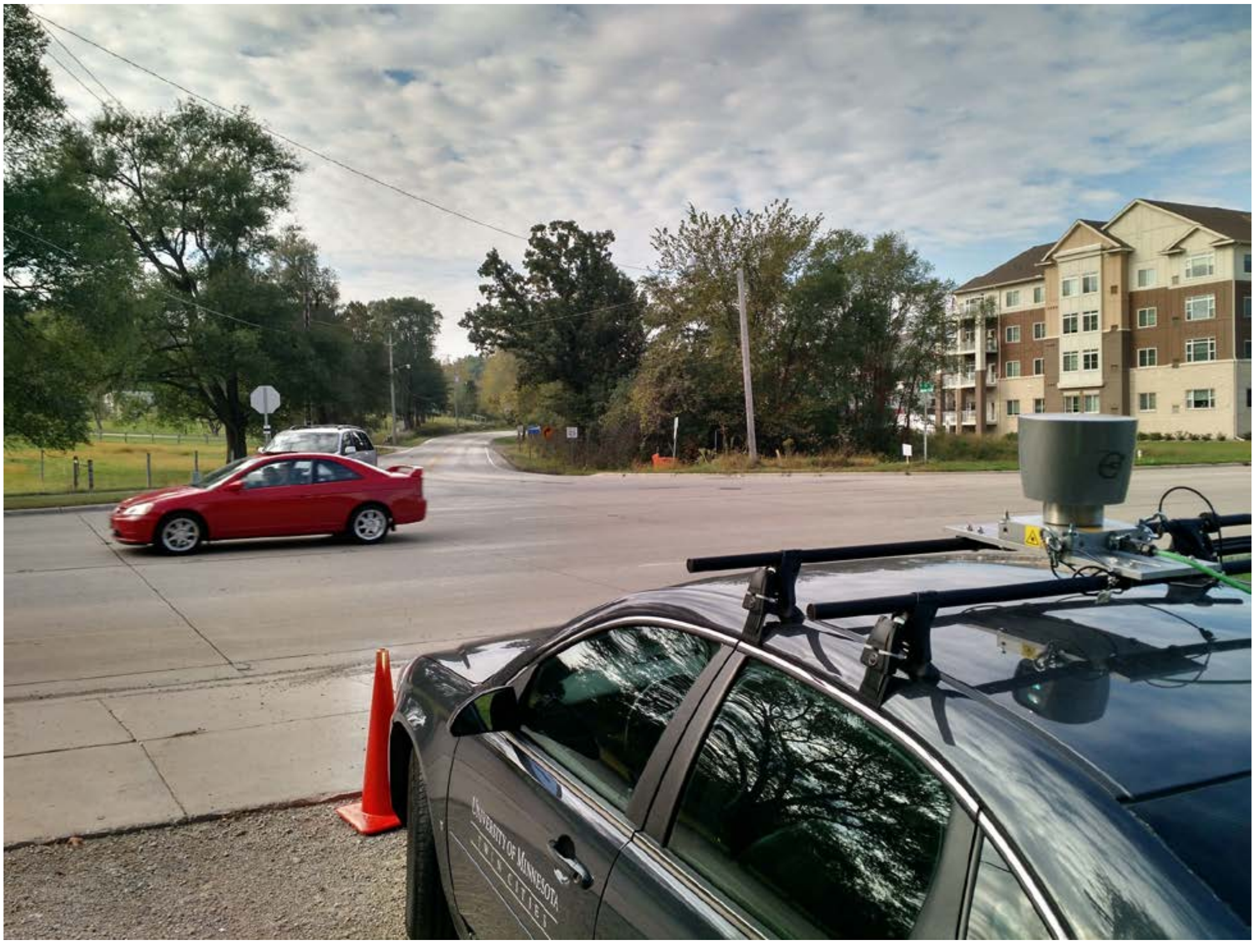

Figure 3.5 Research vehicle stationed at Rochester intersection.

The second data collection effort was in Saint Peter, MN at the intersection of Washington Ave and Broadway Ave. Both of these streets are undivided, 2 lane roads with street parking on both sides. This intersection has a 4-way stop and vehicles making right turns to or from the north had a protected right turn lane that bypassed the stop sign. The intersection is located in a residential area of the city. It’s shown on the map in Figure 3.6.

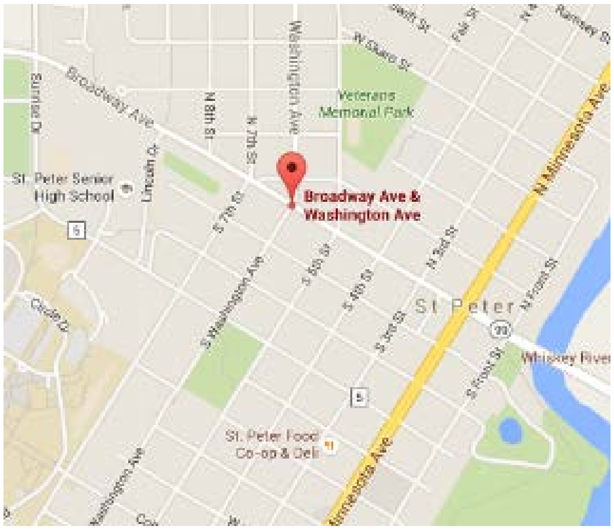

Map data © 2016 Google

Figure 3.6 Location of Saint Peter intersection.

Permission was granted by the city for the research vehicle to be positioned in a painted median on the north side of the intersection. Satellite photography of the intersection is shown in Figure 3.7 which also notes the location where the LIDAR was stationed.

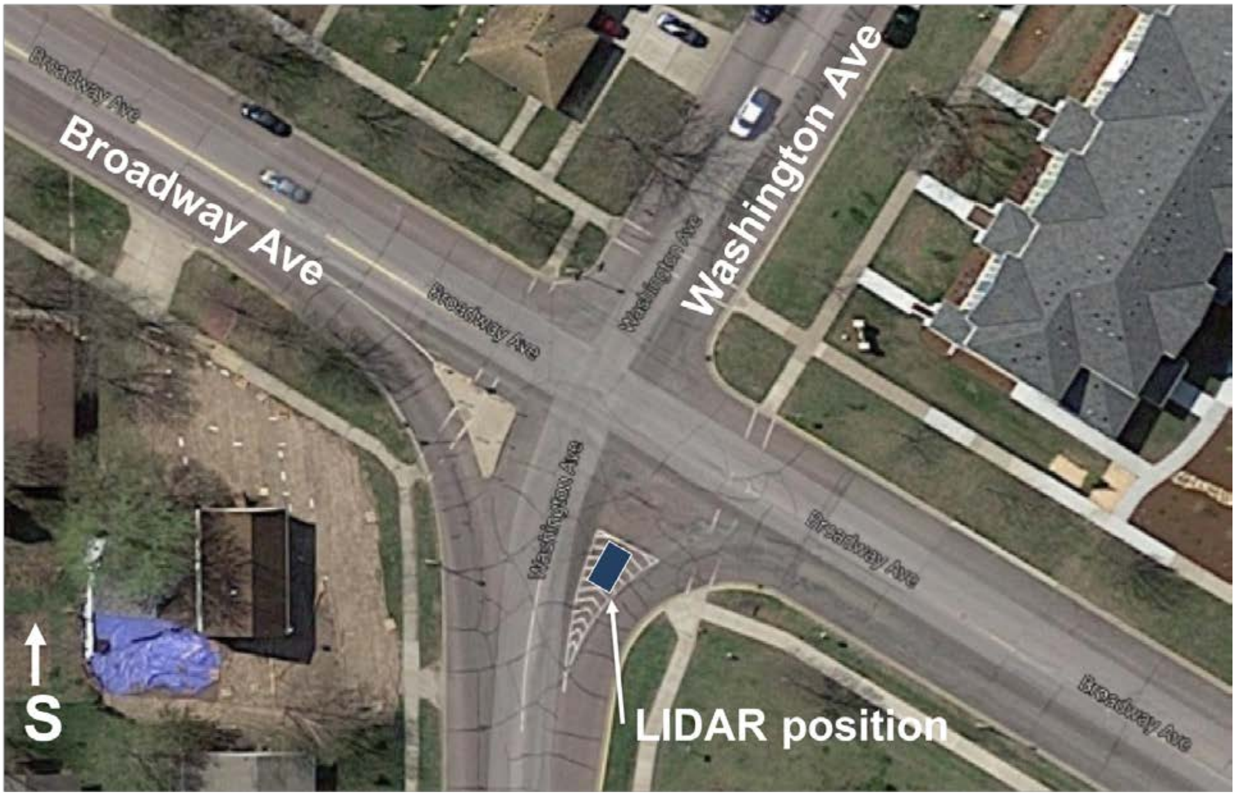

Imagery © 2016 Google, Map data © 2016 Google

Figure 3.7 Satellite photography of Saint Peter intersection.

This intersection was selected for two main reasons. It contrasts the first intersection in that it is a lower volume road located in a residential area. It is controlled with a 4-way stop which would allow for a better investigation of vehicle occlusion. Additionally, by working with City of Saint Peter staff, permission was granted for the vehicle to be stationed in the painted median on the north side of the intersection. This provided a unique opportunity to collect data with the LIDAR as close to the center of the intersection as possible.

Figure 3.8 shows the position of the vehicle within the painted median.

Figure 3.8 Research vehicle stationed at Saint Peter intersection.

3.4 Algorithm Development

The goal of the algorithm development was to create software capable of automatically analyzing the LIDAR data to determine vehicle turn counts. This would be accomplished by first determining the vehicles' trajectories as they move through the intersection and then by aggregating these trajectories by origin/destination pairs to report turn counts.

3.4.1 Software Considerations

The first step in development was determining which (if any) existing software packages should be used. It was decided to use the existing VeloView software as a platform for performing the analysis. This means that the code to perform the processing would be incorporated into the software so that the existing VeloView code would handle decoding the raw sensor output and then new code would be written to handle the post-processing.

The advantages of doing this were that the software already provided a stable interface with the sensor to handle parsing the raw data it generates as well as recording and playing back data files. Additionally, its visualizer could be used to display the output of the post-processing (i.e. tracked vehicle positions).

However, the drawbacks of this approach were that the software wasn’t designed to be used for this type of data processing. It took considerable effort to insert the code used to perform the post-processing and doing so caused the software to be computationally expensive and run slowly. Another challenge was extending the existing drawing/visualization functions to incorporate the additional information to be displayed including automatically detected vehicle positions, trajectories, areas of interest, among other information. Lastly, because data was stored in the raw bit stream format, VeloView would need to re-decode it as it played it back in real-time. This means that to process 10 minutes of saved data, it would take at least 10 minutes because the data would be played back at the speed it was originally collected.

Ultimately, it was determined that the advantages of using VeloView outweighed the drawbacks as it would allow for an efficient prototyping platform noting that future applications would likely require re-writing custom software or investing more time into modifying VeloView.

To perform the point cloud segmentation, the process by which objects in the point cloud are separated and identified as unique from other objects and the ground, two libraries were examined. The first was the Point Cloud Library (PCL) [15] which is a large open source library for 2D and 3D image and point cloud processing. This package would likely be capable of performing many of the desired functions, however using it would require a larger investment to transform the data from the format in which it exists in VeloView to the format required by PCL. The second library examined was OpenCV [16], a computer vision library. This library was better equipped to integrate with the software in an efficient way. For this reason, OpenCV was selected to be used.

3.4.2 Trajectory Tracking Methodology

Once the data was collected at the intersection, the existing VeloView software was used to parse the data and generate the point cloud for each frame (i.e. a single, complete 3D view of the scene created by a single revolution of the LIDAR). This resulted in point cloud data represented by a series of x,y,z coordinates, each corresponding to a single point.

For each intersection, a one-time setup step was performed to identify and define the area of interest. The area of interest is how the geometry of the intersection is represented in the algorithm. It was designed such that all data collected outside of the area was discarded. This allowed the algorithm filter out un-needed data corresponding to the ground, obstructions (signs, lamp posts, other road furniture), and areas outside the intersection. This reduced the number of points that needed to be considered by the algorithm and also filtered the point cloud such that all the remaining points corresponded to vehicles to be tracked.

Figure 3.9 shows a LIDAR scan of the Rochester intersection and Figure 3.10 shows the same view but also draws the area of interest. Similarly, Figure 3.11and Figure 3.12 show similar information for the Saint Peter intersection. Note that although the areas of interest appear to have gaps between them, this is only a graphical issue with the software and they are in fact contiguous.

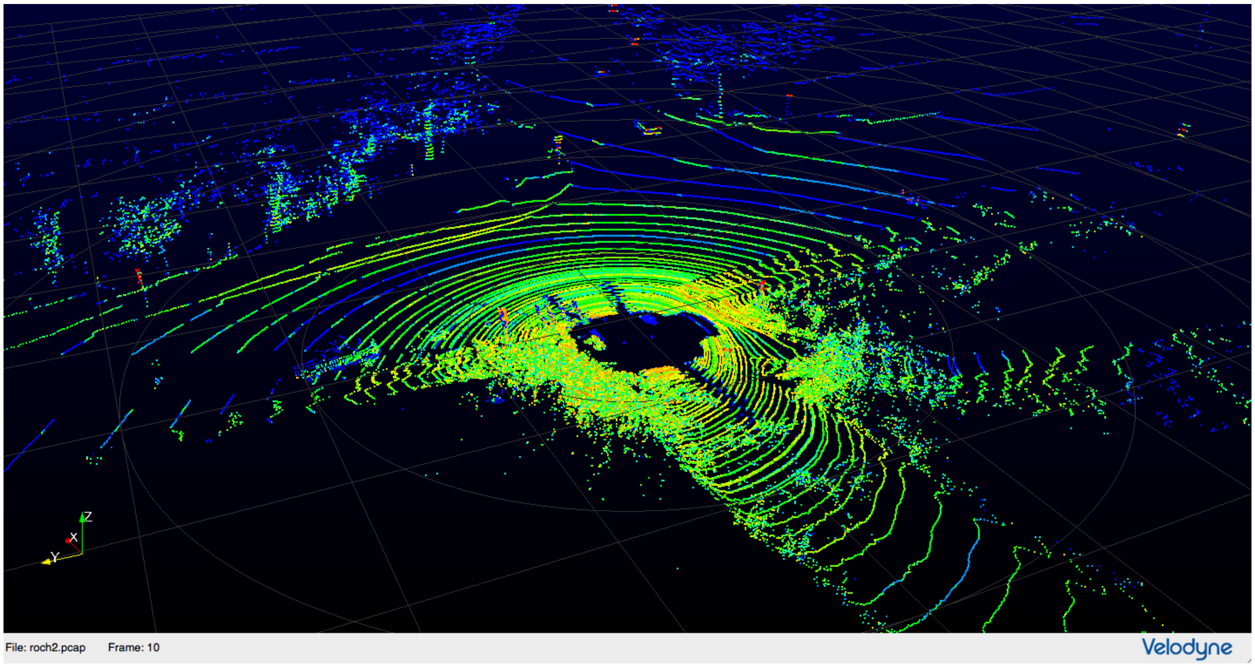

Figure 3.9 Point cloud collected at Rochester intersection.

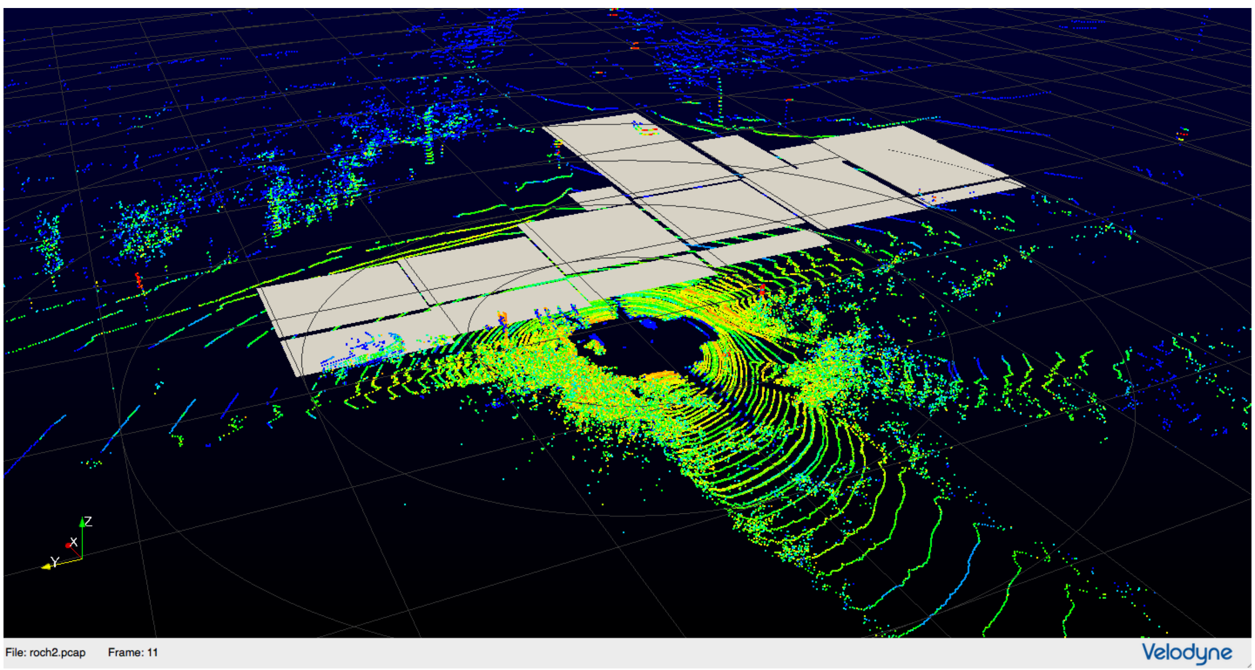

Figure 3.10 Area of interest for Rochester intersection.

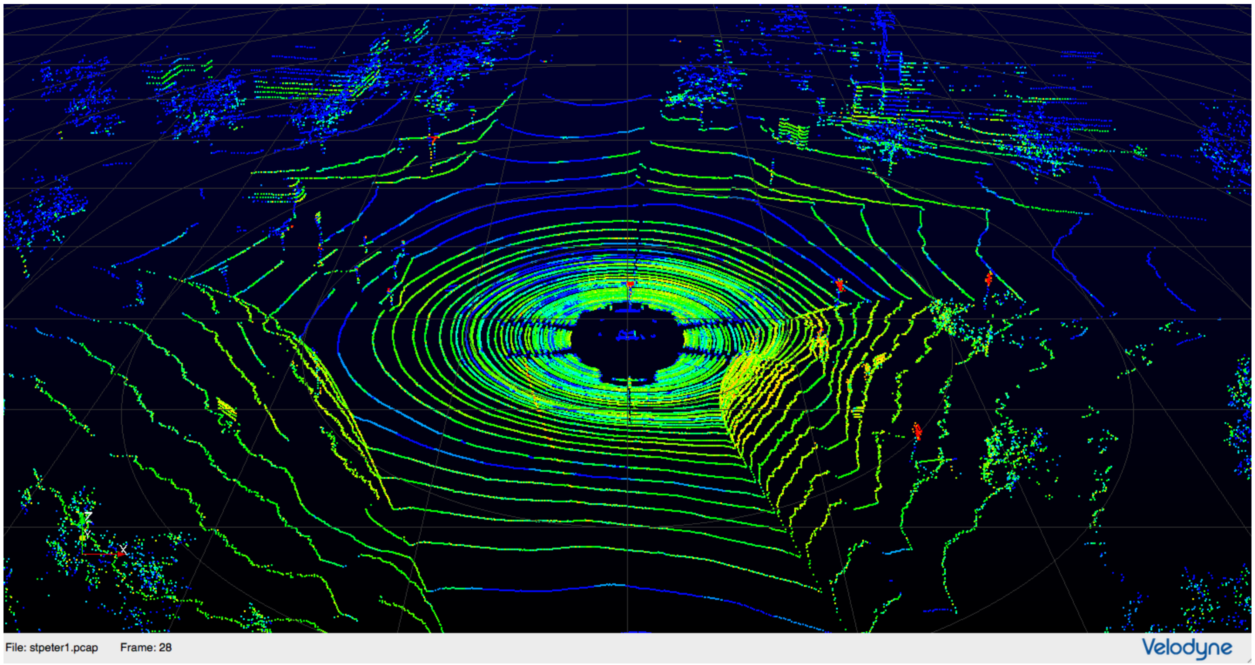

Figure 3.11 Point cloud collected at Saint Peter intersection.

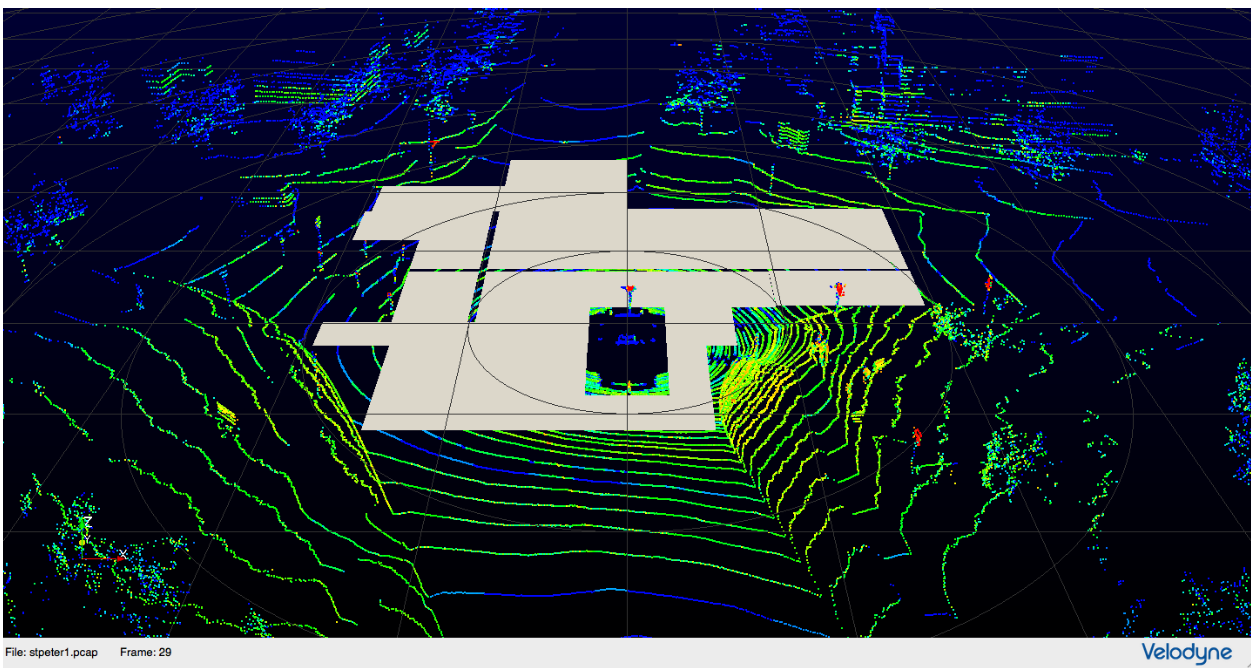

Figure 3.12 Area of interest for Saint Peter intersection.

Note that the areas of interest define not only the area considered to be the intersection horizontally, but they also have an associated height to filter out points corresponding to the ground. Limitations in the software only allowed for the area of interest to be constructed of rectangles whose sides were parallel to the x and y axes of the LIDAR's coordinate frame.

Additionally, entry and exit zones are defined which are used later in the algorithm. These zones correspond to areas where a vehicle is expected to appear or disappear as it enters or leaves the area of interest. An entry and exit zone are created for each lane that leads into or from the intersection, respectively.

With the areas of interest set for the intersections the algorithm is able to use this information as it processes the LIDAR data. The algorithm first removes all height information from the point cloud. That is to say, it only considers the intersection from the top down. Next, it reduces the intersection to a grid of bins, each 0.5 meters by 0.5 meters. For each bin, it counts the number of points bounded by the edges of the bin and uses that to determine whether or not that bin is occupied or empty. Bins with more than 5 points in them are considered occupied and those with fewer points are considered empty. This step is necessary to help filter out noise, only marking bins as occupied if they have enough points in them.

This results in an array of bins that are either occupied or empty. The algorithm then uses OpenCV to connect adjacent bins so that each contiguous set of bins (i.e. bins that are adjacent and bordered by empty bins) are identified and their position is recorded. This is the main step in determining the position of the vehicles within the scene. Here, each contiguous set of bins is assumed to be a vehicle and then its position is saved for the next step. This process is completed for each sensor frame of data.

The next step is to link vehicle positions between frames. This matches vehicles observed in one frame to that same vehicle as observed in the next frame. This allows for vehicles to be identified as they move through the intersection and as they are observed in each sensor frame. This is done by calculating the distance between all vehicles seen in two consecutive frames. Then, if a current vehicle is located within 2 meters of a vehicle in the previous frame, it is assumed to be the same vehicle. Otherwise, it's assumed to be a new vehicle. The result of this step is a set of vehicle trajectories that each correspond to a single, unique vehicle as it moves through the intersection.

The last step is to aggregate this information into vehicle turn counts. This is done by examining each trajectory and recording in which entry zone it was first seen and in which exit zone it was last seen. Additionally, if a vehicle appears of disappears outside of an entry/exit zone, the algorithm looks to link it with a corresponding vehicle that is also lost or acquired in the same location near the same time. This process stitches together incomplete trajectories when a vehicle is temporarily lost in the intersection due to occlusion or other issues. This then generates an origin/destination pair which can be used to characterize the vehicle's movement through the intersection.

3.5 Analysis

A brief analysis was conducted to better understand the performance of the algorithm. It is noted that this evaluation was not designed to be comprehensive but rather to provide insight into the system's performance. The method for analyzing the algorithm was to compare the turn counts automatically generated by the software to counts collected by a human watching the playback of the point cloud. This process was conducted for a single 10 minute collection file from each intersection.

Due to restrictions of the playback software, the algorithm could not process the data any faster than the speed at which it was originally collected. Therefore, processing the 10 minute file took 10 minutes to complete. However, if custom software was developed or additional modifications were made to the playback function, it is likely that this process could have been faster. The processing was capable of running unsupervised, requiring no human intervention once started. This means that a great deal of data could be analyzed in batches for example overnight or in parallel on multiple computers with only a single operator.

The turn counts generated by the algorithm were compared to counts conducted by a human who watched the playback of the LIDAR and hand counted vehicle turns. This process took roughly 20-30 minutes for a single 10 minute file. Additionally, this was effort intensive such that constant attention was required.

The results are summarized in a grid that identifies the number of vehicles entering from each possible direction and exiting from each possible direction. The number on top marked "H" are the numbers collected by the human count. The numbers on the bottom marked "A" are the numbers generated by the algorithm.

The results for the Rochester intersection are shown in Table 3.1.

Table 3.1 Comparison of human vs algorithm turn counts for the Rochester intersection.

	To East	To West	To South
From East	H – 0 A – 0	H – 76 A – 78	H – 12 A – 7
From West	H – 72 A – 75	H – 0 A – 0	H – 12 A – 1
From South	H – 12 A – 1	H – 14 A – 9	H – 0 A – 0

Figure 3.13 shows satellite imagery of the intersection explicitly labeling the directions used in the table.

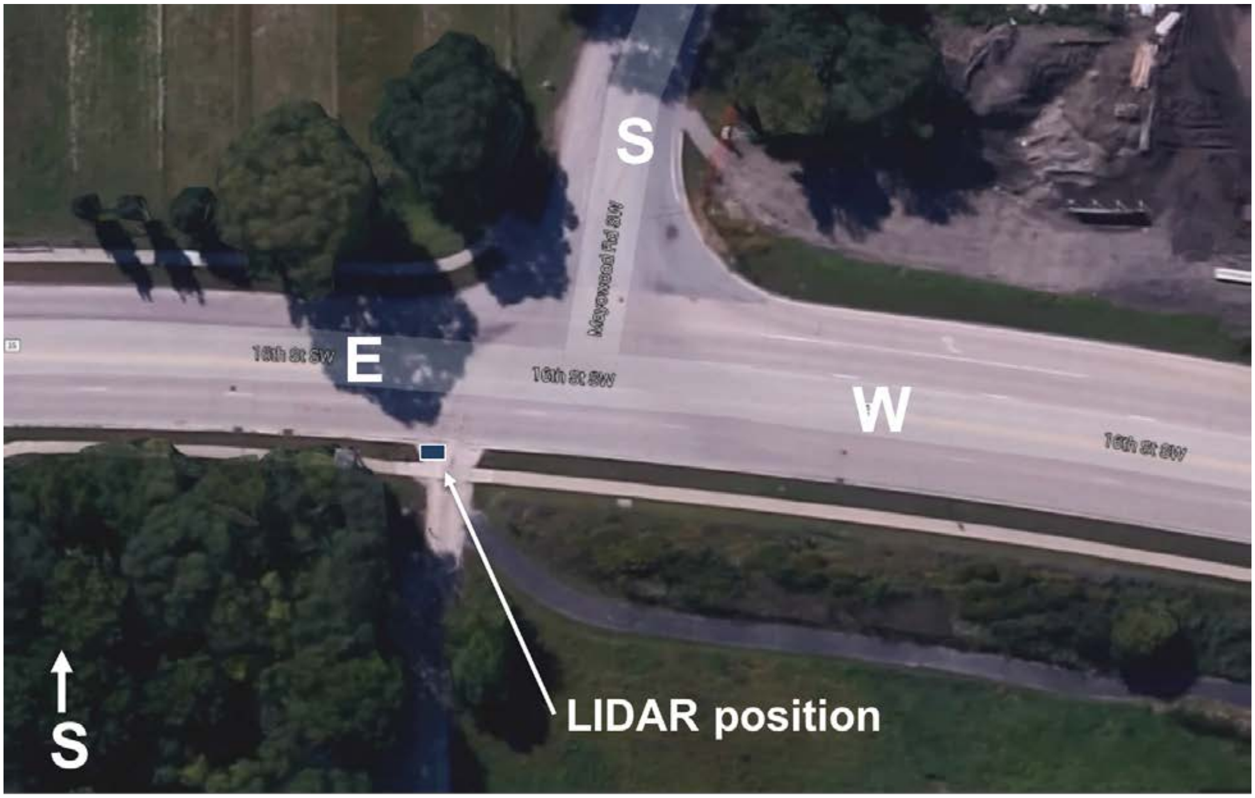

Imagery © 2016 Google, Map data © 2016 Google

Figure 3.13 Rochester intersection diagram labeling approach and exit directions.

The data shows that generally, the algorithm was capable of correctly identifying vehicles moving along the mainline road from east to west or west to east. The most difficult type of turn to capture was when vehicles were making right turns either onto or from the side road. This is most likely due to the distance at which these vehicles were from the LIDAR. Vehicles further away from the LIDAR are represented by fewer points in the point cloud so they become more difficult for the algorithm to detect. Specifically, for traffic making a right turn onto the side road, these vehicles are 5 lanes away from the LIDAR.

For vehicles making right hand turns onto the mainline from the side road (from south to east), a frequent issue was vehicle occlusion. If a vehicle was waiting to turn left from the mainline to the side road (from east to south), it would block the LIDAR's view of vehicles behind it.

The results for the Rochester intersection are shown in Table 3.2.

Table 3.2 Comparison of human vs algorithm turn counts for the Saint Peter intersection

	To South	To North	To East	To West
From South	H – 0 A – 0	H – 20 A – 11	H – 3 A – 0	H – 3 A – 3
From North	H – 12 A – 10	H – 0 A – 0	H – 15 A – 11	H – 9 A – 9
From East	H – 9 A – 8	H – 13 A – 9	H – 0 A – 0	H – 25 A – 25
From West	H – 1 A – 1	H – 0 A – 0	H – 26 A – 24	H – 0 A – 0

Figure 3.14 shows satellite imagery of the intersection explicitly labeling the directions used in the table.

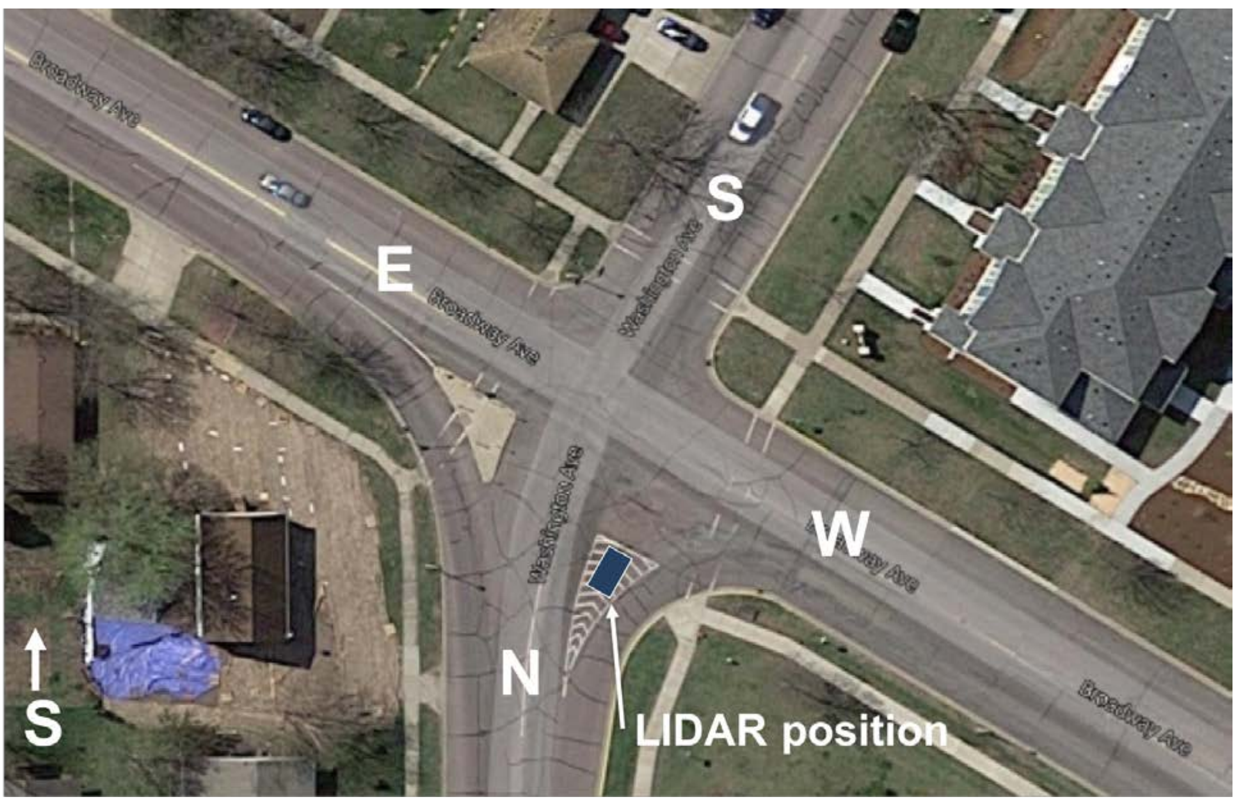

Imagery © 2016 Google, Map data © 2016 Google

Figure 3.14 Saint Peter intersection diagram labeling approach and exit directions.

The Saint Peter intersection had similar results to the Rochester intersection. Vehicles passing across the front of the vehicle, not making turns (from east to west, from west to east) were generally captured accurately. The major discrepancies between the human count and the algorithm count here occurred for vehicles traveling from the south going straight to the north and for vehicles traveling from the east turning right to the north. This intersection was much smaller so range wasn't a significant issue. However, vehicle occlusion was a major issue due in part to the placement of the LIDAR. The research vehicle was positioned very close to the spot where a vehicle coming from the north would stop while waiting for their turn to proceed. This caused large shadows behind these stopped vehicles that would frequently occlude the vehicle movements behind them.

3.6 Discussion

Implementing this sample application provided a great deal of information about both the Velodyne HDL-64E LIDAR, the process of developing LIDAR analysis software, and advantages and challenges associated with using this technology for this application.

3.6.1 System Advantages and Limitations

The major advantage of using a stationary LIDAR scanner to collect intersection data is that with appropriate software, it is possible to collect a great deal of information without traditional, human-based counting methods. Deployment is easy, especially if the sensor is mounted on a vehicle that can be parked near the intersection. The ease of deployment also enables options for limited deployments for example to collect an afternoon of data to characterize traffic due to an event where may otherwise not be cost effective to deploy infrastructure based tube counters. Lastly, the application implemented here was to produce turn counts, but this technology could also provide additional information. For example, it may be advantageous to collect data about not only vehicle turn counts, but their complete trajectories through an intersection. This could allow for the identification of potentially hazardous interactions (e.g. near-misses) between vehicles and other vehicles, bikes, and pedestrians.

The system, as implemented here, also has some limitations. As discussed above, object occlusion can be problematic depending on the placement of the LIDAR. Care must be taken when designing a deployment plan to ensure that the LIDAR will have line of sight to as much of an intersection as possible. If there are particular vehicle maneuvers or portions of the intersection that are of high interest, the LIDAR should be placed such that other vehicles won't block those areas when stopped at a stop line or while waiting to turn.

The methods used to identify vehicles within a point cloud were challenged when observing vehicles further away from the sensor. This is because objects further away from the sensor are represented by fewer points in the point cloud. The way the algorithm was implemented here, the effective range of the sensor was roughly 20 to 40 meters. This does not mean that it is impossible for any software coupled with this LIDAR to detect vehicles beyond this range, but rather as implemented, this version of the post processing software did not reliably detect vehicles observed beyond roughly 40 meters.

Lastly, the algorithm was not designed to automatically detect intersection geometry. An area of interest had to be defined in order for the algorithm to understand the position of the intersection in the LIDAR coordinate frame and also to filter out point cloud points representing the ground, signs, and other stationary, non-vehicle objects.

3.6.2 Future Work and System Improvements

Based on these findings, a number of future improvements were identified to expand the capabilities of the system. First, alternatives to the VeloView software should be examined including making significant modifications to the software or acquiring custom software. Doing this could allow for more efficient post processing and better visualization tools.

Additional development could be performed to enhance the vehicle recognition algorithm to provide higher accuracy and more detailed trajectory information. For example, it may be advantageous to identify time or space headways between vehicles in order to detect potentially dangerous interactions. The detection process could be enhanced to provide not only vehicle presence but also vehicle type (i.e. bus, car, bike, pedestrian, etc.), which would allow for a more complete picture of an intersection's use.

Lastly, the system could be re-designed to use multiple LIDAR sensors. For example, two LIDAR scanners could be stationed at opposite corners of the intersection. Then, their data could be combined and stitched together to create a single, complete point cloud for the intersection. This would mitigate many of the issues associated with the effective system range as well as object occlusion.

3.6.3 Conclusions

The goal of this work was to investigate the use of 3D LIDAR for traffic applications by implementing a sample application capable of reporting turn counts for vehicles moving through an intersection. This work identified advantages and challenges associated with using LIDAR technology this way. It also provided information about how best to address issues with the system for future applications.

CHAPTER 4: CONCLUSIONS

This project had two main goals. The first was to investigate 3D LIDAR and how it may be used in traffic applications by implementing a proof-of-concept demonstration system capable of automatically producing turn counts for vehicles traveling through an intersection. The second was the development of the workshop material including the discussion of the sample application.

4.1 Sample Application

The sample application development yielded valuable information about the strengths and limitations of such a system and provided insight into the identification of future directions to take the application. In short, this application showed that LIDAR might be a suitable tool for collecting traffic data. Future development work would need to be performed to increase system accuracy, mitigate occlusion issues, and implement new features.

Limitations encountered while working with the existing VeloView package highlighted the importance of good software. When working with these sensors, having high quality sensors alone is not sufficient. They must be coupled with effective software capable of performing the required processing tasks. Furthermore, working with software that readily facilitates feature expansion is very valuable.

Before additional development is started however, the most effective next step would be to meet with traffic professionals to solicit additional information about their data needs and how best to accomplish these using LIDAR technology. Feedback about this application during the workshops provided viewpoints not originally considered by the research team. For example, some attendees said that using the technology to get vehicle counts, speeds, and headways along a road segment would also be useful information.

Another consideration raised was that the upfront costs associated with obtaining the hardware used for the sample application could be prohibitive. The discussion included options to mitigate this including a cost-sharing program where a small number of units were purchased for agencies throughout the state to use. Traditional counting methods are already budgeted for and their costs, deployment plans, and accuracies are already well understood such that there may be an institutional inertia associated with fully adopting newer technologies such as this.

Lastly, it was noted that this application provided a proof of concept and although it was clear there were many ways in which it could be improved, it is not yet ready for wide-scale deployment. Reliability is a key factor that would need to be established before such a system could see non-research deployments. However, as the system becomes more refined, more robust, and easier to use, it was hypothesized that due to its advantages, LIDAR-based traffic observation systems may become a valuable tool for traffic professionals.

4.2 Workshop Material

The workshop curriculum was created to provide traffic and GIS professionals with useful information about LIDAR regardless of their existing familiarity with the technology. This was accomplished by starting with the basic principles including how LIDAR works and how to evaluate system specifications when comparing different hardware. In some cases, participants who were familiar with LIDAR in general, still reported finding this information helpful as it solidified their own understanding of the technology.

Next, a number of LIDAR systems from differing industry and application segments were discussed. Even though the focus was on mobile LIDAR, information was provided about aerial and stationary LIDAR. Feedback from participants during the workshops confirmed that these were the types of LIDAR they were most familiar with generally due to their existing knowledge of digital elevation maps and LIDAR for surveying. The goal of this information was to help convey the wide range of hardware available varying greatly in both price and features.

The other major component of the material was the presentation of two applications. The first was about a system designed to create high accuracy maps of roadside features [1] and the second was about the sample application developed for this project. These project descriptions represented an in-depth look at not only the LIDAR hardware in specific applications, but also the processes by which a system using LIDAR hardware is designed and implemented.

The most engaging portion of the workshop was the live demonstration of the Velodyne HDL-64E LIDAR. This was consistently reported by participants to be their favorite portion of the workshop. Although the information presented was useful, seeing the unit in person and viewing the live visualization provided a more qualitative and experiential understanding of the hardware.

4.3 Final Conclusions

This project was designed to provide state, county, and local transportation and GIS professionals with accurate, current, and applicable information about LIDAR systems. To accomplish this, existing LIDAR knowledge was combined with market survey research as well as with new information gathered through the process of creating a sample application. This knowledge was aggregated and used to create a workshop that was informative and well received among participants.

REFERENCES

[1] Davis and Donath. *Development of a Sensor Platform for Roadway Mapping: Part A - Road Centerline and Asset Management.* Minneapolis: Center for Transportation Studies, 2014. Accessed March 2016. http://www.cts.umn.edu/Publications/ResearchReports/reportdetail.html?id=2376

[2] Velodyne LIDAR. “HDL-64E.” Accessed March 2016. http://velodynelidar.com/hdl-64e.html

[3] Sick. “Bulkscan® LMS511.” Accessed March 2016. https://www.sick.com/de/en/product-portfolio/fluid-sensors/flow-sensors/bulkscan-lms511/c/g253553

[4] Trimble. “Trimble AX60.” Accessed March 2016. http://www.trimble.com/Imaging/AX60.aspx

[5] Teledyne Optech. “Optech Orion.” Accessed March 2016. http://www.teledyneoptech.com/index.php/product/orion-altm/

[6] Riegl. “Riegl VZ-2000.” Accessed march 2016. http://www.riegl.com/nc/products/terrestrial-scanning/produktdetail/product/scanner/45/

[7] Teledyne Optech. “Optech ILRIS Terrestrial Laser Scanner.” Accessed March 2016. http://www.teledyneoptech.com/index.php/product/optech-ilris/

[8] PulsedLight. “LIDAR-LITE v2 ‘Blue Label’.” Accessed March 2016. http://pulsedlight3d.com/products/lidar-lite-v2-blue-label.html

[9] Paul Wallich. “Build an Early-Warning System for Your Bike.” *IEEE Spectrum*, June 25, 2015. Accessed March 2016. http://spectrum.ieee.org/geek-life/hands-on/build-an-earlywarning-system-for-your-bike

[10] AutonomouStuff. “ibeo ScaLa.” Accessed March 2016. http://www.autonomoustuff.com/ibeo-scala.html

[11] Riegl. “Riegl VQ-450.” Accessed March 2016. http://www.riegl.com/nc/products/mobile-scanning/produktdetail/product/scanner/31/

[12] Riegl. “Riegl VMX-450.” Accessed March 2016. http://www.riegl.com/nc/products/mobile-scanning/produktdetail/product/scannersystem/10/

[13] Velodyne. “HDL-32E.” Accessed March 2016. http://velodynelidar.com/hdl-32e.html

[14] Velodyne “Puck™ VLP-16.” Accessed March 2016. http://velodynelidar.com/vlp-16.html

[15] Open Perception. "PCL – Point Cloud Library." Accessed April 2016. http://pointclouds.org/

[16] Itseez. "OpenCV." Accessed April 2016. http://opencv.org