

BCOE WHITEPAPER

Smart Transportation Systems

Introduction

The following report was produced based on a grant from the University of New Hampshire's Broadband Center of Excellence (BCoE) to the UNH community seeking ideas on the use of smart sensors to benefit rural or disadvantaged communities. BCoE provides unbiased information and demonstrations of broadband for the development of innovative network application experiments in the education, health, public safety and economic development sectors to improve citizens' quality of life. BCoE seeks to attain affordable broadband available to all people around the world.

Participants

PI:

Nicholas Kirsch, Electrical and Computer Engineering

Staff:

Christina Dube, Connectivity Research Center

Co-PIs:

Andrew Kun, Electrical and Computer Engineering

Robert McGrath, Health Management and Policy

Prashant Mittal, Health Management and Policy

Students:

Cameron Devine, Connectivity Research Center

Maxwell McNally, Connectivity Research Center

Executive Summary

Vehicular traffic of all kinds is vital in semi-rural communities such as Durham, NH. While our community does not have the traffic congestion of larger cities, our vibrant pedestrian contingent, cars, buses, and maintenance vehicles are sometimes at odds. The challenges range from pedestrians in un-signalized crosswalks between classes to the lack of parking spaces during maximum demand^l. These problems are universal throughout the country and scale to much larger cities.

Smart cities are communities that utilize information from sensors and sensing systems to make services more efficient or improve sustainability. These systems are essentially a cyber-physical system that acts upon gathered information that is transmitted through a wired or wireless networkⁱⁱ. In this context, we are interested in deploying a network of sensor systems throughout UNH and Durham to improve transportation in our community.

The **goal** of this project is to have a positive impact on the community by improving awareness of parking, decreasing traffic congestion caused by pedestrians and researching novel methods to encourage the use of alternative forms of transportation. The **innovation** of this project is the synthesis of data analytics from networked sensors and new forms of displaying information to community members.

There were several outcomes of this project and we achieved many of our goals:

- Developed and deployed two sensor platforms (parking and pedestrian); these platforms included sensor architecture and devices, connectivity to the backend, and necessary servers.
- Performed a range of analyses on the collected data that demonstrates the efficacy of making these measurements. Certain limitations of the sensors did not enable us to collect sufficient data to build a complete UNH transportation model.
- Created a platform for displaying transportation information to a motorist through augmented reality and to understand the impact of using these tools while driving.
- Brought a multi-disciplinary group of UNH faculty together to work on a Smart Cities project that lead to the submission of a federally funded research proposal and connections to local transportation companies and DOTs.

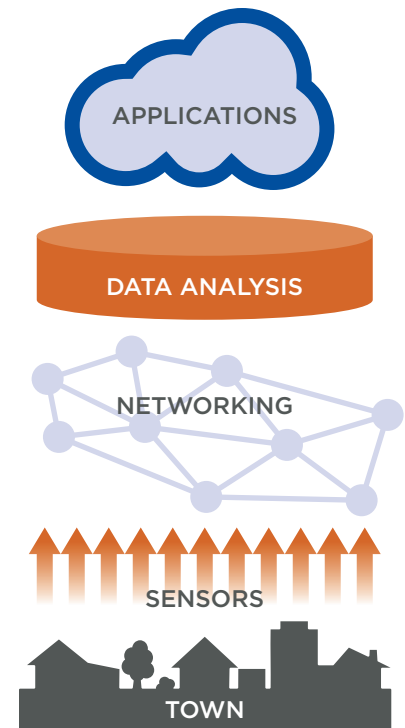


figure 1: Smart Cities tiers

Sensor Platforms

The first objective of our project was to develop two sensing platforms for measuring parking availability in one of the University lots and pedestrian traffic in key crosswalks. The proposed parking system has a set of solar powered wireless sensors to count the number of spaces available in a lot and to communicate with a server. The pedestrian sensor monitors mobile phone emissions as pedestrians walk by the sensors and will similarly store the data.

Parking Sensor

An ideal solution would require little to no maintenance for the life of the project and would be able to be deployed remotely, with the data accessible in real time, e.g. burying sensors under every parking space. However, this effort was outside of the scope of the project. Instead, we developed a temporary sensor platform. When considering the amount and frequency of the data, LoRa wireless technology was the best fit since it provides a low power solution that could be accessed at a long range. The decision to use LoRa, then drove the processor selection, which was MultiTech's mdotⁱⁱⁱ processor.

The main requirement for the sensor platform was the ability to detect a vehicle passing through the entrance of the parking lot. There were a variety of sensors that met this requirement. Magnetic, infrared, LIDAR and ultrasonic sensors were explored, and ultrasonic sensors were deployed.

Initially magnetic sensors seemed to be the logical choice, however, to detect a vehicle, the magnetic signature (waveform) had to be analyzed and compared to a "known" vehicle magnetic signature.

Without a user actively reading the data, this analysis required a DSP filter. The DSP filter would have required more power for processing than could be afforded on the project.

Infrared tests showed promise when detecting objects in the lab, however infrared was not chosen because the sun's infrared spectrum emission confused the sensor because it is in the same wavelength as the sensor. LIDAR is more expensive than our solution warranted.

Ultrasonic sensors could accurately detect vehicles up to 10 feet from the sensor platform. Given that the entrances to the parking lots are large areas of space with cars turning from various directions, this was the most ideal sensor for the project.



Figure 2: Original parking sensor form factor as deployed in Parking Lot B

Block Diagram




Figure 3: Block Diagram of Parking Sensor

Each sensor module contains a single mDot as the processor handling sensor input and network connectivity. Two Ultrasonic range sensors are connected to the mDot, one monitoring the inner-side of the parking lot entrance and another monitoring the external side. These transmit ultrasonic pulses and then receive the pulse back. Upon reception, they monitor the time of flight of the waveform. Knowing the speed of sound, the distance from the sensor to the car can be measured. If the distance measured is within a certain threshold then that sensor has detected a car. If the other sensor senses a car in the same cycle, the direction of the car can be determined.




Figure 4: Processor module wired to sensors

The sensors used were HC-SR04 Ultrasonic sensors. They have four pins, a power and ground for powering the device, and a trigger and echo pins. When the trigger pin is sent high the sensor sends out an ultrasonic pulse that is then received by the sensor. The echo pin returns a high pulse that corresponds to an object being in the field of view of the sensor. The time of flight (difference between send and receive) of the ultrasonic pulse is then used to calculate distance. The module counts vehicles passing by in 15-minute intervals.

People and other passing objects are purposely ignored by distance thresholds and time constraints on the sensed object, but this is not perfect. When 15 minutes passes the device uplinks to the LoRa network and transmits a frame with data indicating “Cars In” and “Cars Out”. The data is then pushed from the Senet network to a server on the UNH campus for further analysis.

In addition to sending data up to the network, a downlink can be sent to the device through the network immediately after an uplink to either reset the device or to have it sleep. The devices were often told to sleep overnight to save additional power.

These sensors were deployed at the entrances to the A-Lot and B-Lot parking lots on the UNH campus. For one entrance of B-Lot, due to the large width of the entrance, two sensor systems were deployed, one on each side of the entrance. Another similarly sized entrance on the same side of B-Lot could not be instrumented on both sides due to the proximity of a fire hydrant, and the inability to secure the device in place.



Figure 5: New form factor of the parking sensor as deployed in Parking Lot A

Location of Parking lot sensors (in yellow)



Figure 6: Three sensors deployed in Parking Lot A at the entrances



Figure 7: Four sensors deployed in Parking Lot B at the entrances

Pedestrian Sensor

Detecting people was a challenging, yet vital, goal in order to make transportation more efficient and safe. The ability to do this would allow better traffic management and allow commuters to better plan their drive around densely packed areas. Current solutions for this problem are expensive, complex and use a lot of power.

The current solutions include computer vision, analyzing WiFi traffic or analyzing cell phone signals. These can be extremely complex and for this project, consume too much power. Many highway departments now use Bluetooth signals to analyze car travel times on roadways. Instead of travel times, counting the amount of Bluetooth devices can be related to the amount of people in the area. Bluetooth has a small range of around 10 meters so the count would be relative to that small area.

One problem with Bluetooth is that not all people carry a Bluetooth device and even when they do it is not always activated and detectable. Because of this, an experiment had to be conducted to create a probability model of how many Bluetooth devices there are per person. These experiments allowed us to correlate the number of devices and the population in a small area and then accurately predict the population using the data we previously recorded.

A Raspberry Pi 3 was used as the processor module. The Raspberry Pi has a Bluetooth module on it with the BlueZ Bluetooth stack that allows the user to control the module. A python script was written to scan and record the MAC addresses it could detect and get a count of Bluetooth devices by counting the number of unique addresses. Each MAC address is guaranteed to be linked to one Bluetooth device.

Block Diagram

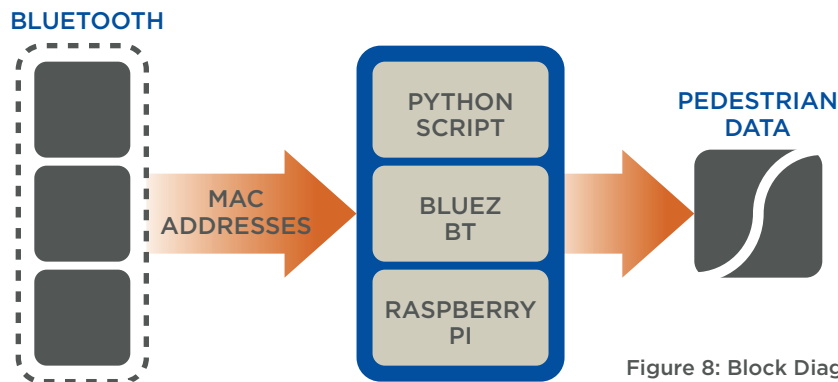


Figure 8: Block Diagram of Parking Sensor

The experiments were conducted in Durham in front of Thompson Hall at the intersection of Garrison Ave. and Main St. Every sample taken would count Bluetooth devices in the area for 30 seconds. During those 30 seconds, the amount of people walking through the area were being counted. At the end of the experiment, the number of devices was compared to the actual amount of people so in the future, the prediction of the amount of people in an area compared to the number of devices can be made with high confidence.

In the future, these devices can be upgraded to move away from an entire Raspberry Pi using a lower power Bluetooth module instead, with LoRaWAN to provide access to a central server at low power and cost. This data can be used to give real time predictions of population to users in Durham.

Data Analysis


The purpose of the data analysis is to analyze the base-line and sensor data from the first objective to reduce the stress level for drivers and pedestrians, but also to help direct the traffic around the campus in a safer and more efficient way. In the future, the broadcast of such information through an app or signage requires complete understanding of the current traffic situation, number of parking spots available all over the campus and highly accurate predictive models that constantly modify their predictions of availability of spots.

Parking Sensors

The csv file containing time series data points for all days between the third week of January through the second week of March were loaded into memory in a manner that allows for detailed analysis for B lot. Measurements were also made for A lot. This process simply renames columns with intuitive names, creates a date formatted column, numerical weekday column (i.e. Monday 0, ..., Sunday = 6), and parsed time stamps into hourly columns. The data were 15 minute counts of cars entering in and exiting the lot. Sensors had a systematic tendency of overestimating cars exiting the lot and underestimating those coming in the lot. We calibrated sensors for various measurement errors several times, and there still existed estimation anomalies.

Our sensor platform was not completely stable and led to anomalies that made it difficult to create a complete transportation model based on the time when each lot was at capacity. Several methods of detecting anomalies were employed. The first method tried was an implementation of an Isolation Forest^{iv} algorithm. The details of this algorithm are not discussed as it is beyond the scope of this report. The second method implemented was a rule based anomaly identifier. The rules rely on grouped averages and standard deviations to flag a record as anomalous or not. The function is built to flag records as anomalous with respect to either cars in or cars out. A record is flagged as anomalous if its number of cars in or out is greater than n standard deviations above the mean of cars in or out grouped by hour and quarter of the hour.

The third method of anomaly detection was an “out of the ballpark” range of number of cars in and out. For instance, there were cases where number of cars exiting in an hour exceeded nine hundred, which is physically impossible and indicates anomalies in sensor detection of vehicles. However, we present a data set to demonstrate the efficacy of our sensor platform by analyzing the data from the sensors. The following graphs depict the number of cars entering and exiting each lot on campus.






Figure 9: 24-hour view of cars entering and exiting Parking Lot A. Each color represents a day during the week of April 24 – April 28


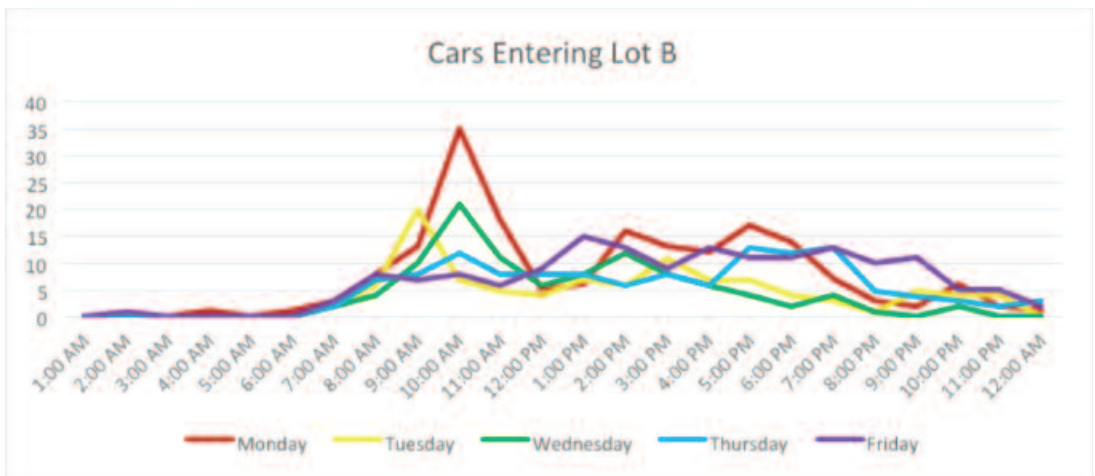


Figure 10: 24-hour view of cars entering and exiting Parking Lot B. Each color represents a day during the week of April 24 – April 28

One interpretation of the data is that A lot rarely reaches capacity and cars arrive in the morning and leave in the evening. However, in B lot cars arrive throughout the day and leave as early as 9AM. This leads us to believe that the lot is full when a car enters, looks for a parking place, is not successful, and then leaves.

Pedestrian Data

The experiments were conducted in Durham in front of Thompson Hall at the intersection of Garrison Ave. and Main St. Every sample taken would count devices in the area for 30 seconds. During those 30 seconds, the amount of people walking through the area was being counted. At the end of the experiment, the number of devices scanned was compared to the actual amount of people so that in the future, the prediction of the amount of people in an area compared to the amount of devices can be made with high confidence.

The raw data over two experiments shows the correlation between the number of devices detected and the number of people counted.

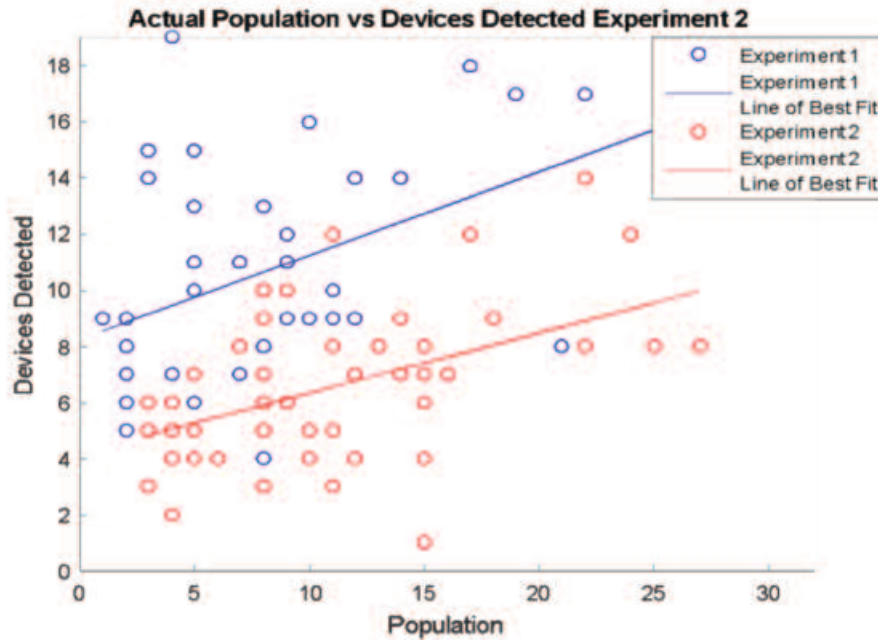


Figure 11: Actual Population vs Devices Detected for each Bluetooth Experiment

| | DEVICES PER PERSON | |
|--------------------------------|--------------------|--------------|
| | EXPERIMENT 1 | EXPERIMENT 2 |
| Mean | 1.86 | 0.74 |
| Median | 1.22 | 0.66 |
| Standard Deviation | 1.61 | 0.41 |
| Correlation Coefficient | 0.49 | 0.48 |

Table 1: Statistical analysis of Devices per Person for each Bluetooth Experiment

The data demonstrate a positive correlation of about 1.5 devices per person. We also analyzed the density of devices per person and found that the distribution is likely Laplacian distributed. However, a larger study is required to create a better model.

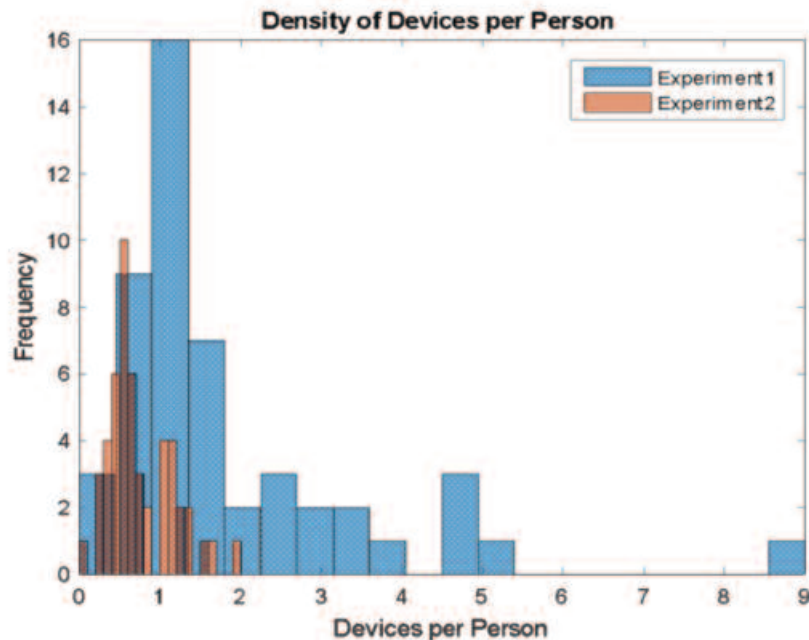


Figure 12: Density of Bluetooth devices per person for each experiment

These experiments proved that Bluetooth can be used to predict the amount of people in an outdoor area while keeping cost and power required low. Creating a strong probability model can make these predictions more accurate. These experiments showed that 1.5 devices can be equal to 1 person or 15 devices will be seen for every 10 people. This is the case 80 percent of the time with a confidence level of 90 percent. More experimentation can be done to get a higher percentage with a greater confidence bound.

Transportation Augmented Reality

Augmented reality holds the promise of providing relevant information to road users in a manner that will allow the road users to remain connected to their surroundings. This connection to their surroundings has multiple benefits. First, for drivers, this approach should allow them to consume information, such as navigation instructions, while they retain sufficient attention on the outside world to safely perform the primary task in the vehicle, which is driving. Second, for passengers (and we will all be passengers in automated vehicles), the connection to their surroundings means that information can be displayed in a context-sensitive manner. Context sensitivity in this case means that information is presented such that it is visible under different road and environmental conditions, and that the information is closely tied to the physical world that it relates to.

To make augmented reality work in a smart transportation system will require significant broadband capabilities. Drivers and passengers alike will want to see rapidly changing visual information that is tied to their physical context. Examples include visual augmentations of the physical environment for gaming purposes, as well as navigation instructions that augment physical landmarks with virtual landmarks. Thus, we expect that augmented reality applications will constitute a significant portion of broadband data consumption in smart transportation systems.

To make augmented reality work, we need to answer questions related to in-vehicle human-computer interaction. What information do drivers and passengers want to consume in vehicles? How should the information be presented? For example, which virtual objects can be used as virtual landmarks for navigation applications? And to answer these questions we need some tools. We feel that one critical tool is to be able to provide accurate data about where users are looking at any given time as they are using augmented reality.

In work at the Human-Computer Interaction Lab at the University of New Hampshire we tackled two of the issues mentioned above. First, we explored the question of how drivers might use augmented reality, and specifically, we wanted to find out if they would use augmented reality to complete video calls with remote conversants. In a driving simulator-based experiment we found that one current implementation of augmented reality — the Microsoft HoloLens augmented reality glasses — might discourage drivers (and possibly also passengers) from looking at the video of a remote conversant (Figure 13). This is because HoloLens has a narrow field of view (about 40 degrees wide), which might require the user to move their head often to see the video.

We are also working on integrating eye tracking with HoloLens, to allow us to precisely measure users' gaze angles both in the virtual world of HoloLens and in the real world that surrounds the user. Preliminary data indicates that we can measure eye gaze location accurately both for real targets (Figure 14) and for virtual targets.



Figure 13: In a driving simulator-based study we explored the use of HoloLens to complete video calls between drivers and remote conversants.




Figure 14: Head-tracking data from HoloLens was integrated with gaze direction data from a wearable eye tracker. This figure demonstrates that this data can be used to accurately measure gaze location in the real world for users wearing HoloLens.

Conclusions

In this smart transportation project, we developed and deployed two complete wireless sensor networks, measured and analyzed two sets of transportation data and created a novel augmented reality platform to display information to motorists. Each of these tasks brought the research group closer together and created a vital data set to apply for future Smart City research projects.

In addition to the team working together as researchers, other valuable connections and relationships were developed that will enable future Smart City work on the UNH campus. These relationships include working with the UNH Transportation and Parking departments, the UNH-Police department and the Town of Durham Police Department.

Many other tangential items were learned and established that could also enable future Smart City work. In the research phase of the sensor development, many processor platforms, sensor types and wireless technologies were explored which will benefit other projects. A LoRa network has been deployed on campus, which includes a server for the sensor data. This network can be utilized for future projects. Most importantly, many lessons were learned regarding the actual deployment, monitoring and maintenance of a remote sensor network in the New Hampshire climate. This knowledge will be critical to future Smart City projects at the University of New Hampshire.

- i Nagatani, Takashi. "The physics of traffic jams." *Reports on progress in physics* 65.9 (2002): 1331.
- ii Jin, Jiong, et al. "An information framework for creating a smart city through internet of things." *Internet of Things Journal*, IEEE 1.2 (2014): 112-121.
- iii <http://multitech.com>
- iv Liu, Fei Tony, Kai Ming Ting, and Zhi-Hua Zhou. "Isolation Forest." 2008 Eighth IEEE International Conference on Data Mining (2008)

© 2017 University of New Hampshire
Broadband Center of Excellence

For more information, visit

www.unhbcoe.org