Technologies for Safe & Efficient Transportation

THE NATIONAL USDOT UNIVERSITY TRANSPORTATION CENTER FOR SAFETY

Carnegie Mellon University | UNIVERSITY of PENNSYLVANIA

Crowdsourced Traffic Calming

FINAL RESEARCH REPORT

Bob Iannucci, PI
Ceferino Gabriel Ramirez, PhD Student
Ervin Teng, PhD Student
Joao Diogo Falcao, PhD Student
Kim Jin, MS Student
Samarth Mathur, MS Student
Anand Sankar, MS Student
Prajwal Prasan, MS Student
Ashutosh Tadkase, MS Student
Nagarjun Srinivasan, MS Student

MAP-21 TSET National (2013-2018)
Grant DTRT-13-GUTC-26
Project No. 14

DISCLAIMER

Crowdsourced Traffic Calming

Abstract

We present our work on creating key elements of a Crowdsourced Traffic Calming system as an enabler for non-specialist citizens to engage in traffic measurement and evaluation in their own communities. We present our design of a roadworthy, low-power, wireless Traffic Sensor; deployment of a low-cost, wide-area sensor network supporting efficient and economical gathering of traffic and other smart city data; a set of browser-based visualization tools for analyzing acquired traffic data; and a set of related studies to optimize the communication and energy performance of sensor within the context of this new network. We discuss the challenges uncovered while deploying this system and hindrances that may be encountered by other city-scale, traffic-related studies.

Keywords: Traffic Calming, Sensor Networks, LoRaWAN

I. Problem

Traffic calming (TC) is an approach to moderating vehicular traffic speeds that relies on the psychological and practical effects of lane narrowing, speed tables, lane deflection, restricted access and similar interventions. Traffic calming has been shown to be effective in reducing accidents, reducing effective speeds, reducing noise from road traffic and reducing the length of waiting time for pedestrians to cross the roadway [1].

The choice of specific traffic calming measures can be situational and is often approached experimentally. For example, some municipalities maintain stores of movable rubber curbs and other devices that can be used to prototype TC interventions, even offering neighborhood groups the opportunity to conduct these experiments themselves. Such approaches can lead to an effective outcome, but the tools for applying quantitative measures for evaluating alternatives all too often are only basic–such as time-limited deployments of pneumatic-hose vehicle sensors or a few people with clipboards to count cars. As a result, the information collected is limited, making quantitative analysis of traffic flows with and without specific calming interventions difficult.

Significant research into wireless device engineering, mesh networking, energy harvesting over the last 15 years has led to the possibility of Smart Cities. Pervasive sensing and wireless networking in the Smart City hold out the very real possibility of enabling very fine-grained, real-time, and ongoing traffic monitoring and quantification, leading to improved metrics for traffic calming interventions. Pervasive sensors, an infrastructure for collecting data, and real-time tools for tracking traffic behavior–and comparing data across different TC techniques–could turn today's ad-hoc practices into large-scale science. Within the vision of the smart city are sensors for all sorts of physical quantities including air quality, pedestrian presence, empty parking spaces, and vehicular traffic. But while the sensors themselves may be inexpensive, the wireless infrastructure that connects them may not be so. Current trial smart city deployments based on cellular networks incur high cost-per-sensor for operation (monthly charges), as well as devices (radios) themselves that are power-hungry, necessitating costly maintenance of battery-powered devices. A top-down approach to outfitting a city with pervasive sensing is in general hard to justify without a direct tie between investment and payoff (such as reduced costs, increased resident satisfaction, or other such metric) [2].

New wireless technologies are emerging that can reduce or eliminate these barriers. LoRaWAN [3], a relatively new low power wide area networking (LP-WAN) scheme that uses Semtech's LoRa [4] modulation technology, shows significant promise in providing the long distance, low-power, and low operating cost wireless interconnection demanded by applications such as smart cities where distances and physical obstacles are significant and the need to conserve power on small, battery-powered sensors is extreme. [5], [6]

This project presents our approach to Crowdsourced Traffic Calming–a system of elements that are simple enough to use as to enable non-specialists to engage in traffic science at levels from neighborhoods through cities. We elected to start from prior work in low-power sensing and build a complete system from sensor to network to visualization and analytics as a system proof-of-concept and to identify gaps in the state-of-the-art that stand between this proof-of-concept and broader deployment. Specifically, we identified the need to explore more sophisticated on-sensor processing, leading to our development of an advanced, low-power sensor for capturing and recording traffic flows. We adapted emerging technologies for low-power, wide-area networking to create a low-cost network supporting wide-area data collection for not only traffic data but also broader smart city data. We also identified the opportunity to re-cast smart city sensing as an open, integrative activity that can engage third parties more readily. In support of this, we created a flexible, adaptable browser-based visualization system for analyzing traffic flows. We present the results of this study and offer our perspective on possible future directions.

II. Approach

Our approach begins with considerations of the principles to which any satisfactory solution should adhere. We then examine the shortcomings of traditional cellular networks as they relate to wide-area traffic sensing and consider the opportunities that are opened with the advent of low-power, wide-area networking. We further consider the benefits and challenges of making sensors and their application simple enough that non-specialists could apply them. These considerations motivate the creation of a new kind of smart city traffic sensor, the characteristics of which we outline. These concepts are also discussed in a companion paper published at the ITS World Congress in Montreal [2].

A. Design Principles for the Crowdsourced Smart City

If we assume that it is desirable to make cities smarter through novel applications of sensing, computing and actuating technologies to enable applications such as quantifiable TC, we do so with some tacit assumptions in mind. The kinds of applications we contemplate depend on (a) the ability to place devices *anywhere appropriate for the purpose* in the smart city; (b) city-friendly approaches to introducing these devices into a city's physical infrastructure; (c) presentation of smart city resources as an *open programming platform* that will attract the same millions of programmers who made the mobile computing revolution what it is today; and (d) a seamless view of devices, network and cloud that make programming approachable. We examine each of these in detail.

For the sake of exploring the challenges of smart cities in general, and traffic analysis in particular, we set aside consideration of in-building environments. This is not to say that smart buildings are not part of the smart city but rather to say that the subject of in-building techniques for sensing and actuating are relatively well-understood [7]. Rather, we choose to focus on the less-well-understood *out-of-building experience*–specifically, traffic and similar outdoor applications.

1) Devices Anywhere: Cellular networks have changed the way we think about network-connectedness. It has only been three decades since Nicholas Negroponte challenged our thinking of television (then over-the-air) and telephone (then largely wired) by asserting they would exchange their delivery modalities (referred to by George Gilder as the *Negroponte Switch*[1]). Now, wireless connectivity for phones is presumed, and this presumption spills forward to the Internet of Things. Filling the world with hundreds of billions of sensor devices can only happen with pervasive wireless networking. Smart cities inherit this assumption if only due to simple economics. The installation cost of a small sensor that connects wirelessly and harvests its own energy is dramatically lower than one that mandates the installation of power or network wires. If it takes a crew of four a day to install the power and/or network wires for a fifty-cent sensor, then all notion of deploying a dense sensor fleet evaporates. Moreover, if each such device has to be periodically serviced (*e.g.,* having its battery replaced every month, having its firmware re-flashed as needed), the per-device per-year maintenance costs will likely render the entire solution infeasible.

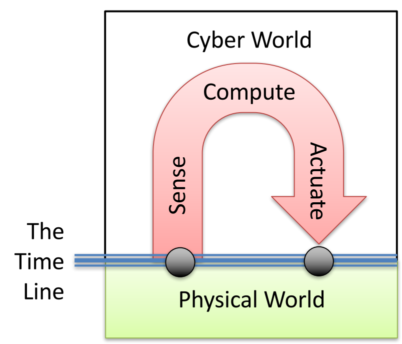

Fig. 1: Simplified representation of information flow in the smart city. Physical quantities (*e.g.,* passing cars) in the smart city are sampled by sensor devices. In so doing, their values are associated with the time-of-reading and passed on for processing. Decisions may result in actuation (*e.g.,* traffic signal changes, application of brakes). In both cases, information crosses the Time Line from the physical world where time has meaning to the cyber world where it is merely an information tag.

Consequences: We must assume pervasive wireless connectivity and, just as importantly, self-poweredness (either harvesting or shipped-with-energy-for-life). Are today's cellular networks appropriate?

2) City-Friendly Integration: Chief among the considerations in creating a smart city network are minimization of cost and maximization of alignment with standing city practices. Creating new procedures for city workers,

[1] https://en.wikipedia.org/wiki/Negroponte_switch

educating installers, and designing new maintenance practices at city scale can result in costs that could overwhelm the perceived benefit.

We consider Pittsburgh, Pennsylvania as a representative mid-sized US city. It owns and maintains 1,031 miles of roads[2]. For our application, we can imagine a dense array of wirelessly-connected sensors, placed every 100 feet or so along these roads to gather high-quality traffic data. This application alone would require installation, recording the position, and optimizing the radio performance of over 500,000 individual devices. We think it unlikely that Pittsburgh or any similar city would employ a small army of specialists for the purpose of geo-locating and optimizing the installation of these devices. We can, however, imagine that the city could install the entire fleet over time if the sensing technology could be integrated into road fixtures that already have established installation and maintenance procedures.

Consequences: Integrating with traditional city infrastructure implies a level of hardened packaging quite unlike in-building applications. Unobtrusive placement (*e.g.,* on the pavement) may lead to poor conditions for wireless antenna performance that will have to be made up by the network. City integration implies that the maintenance intervals for our sensor devices will have to match those of the host infrastructure. Lacking specialist installation, the burden of performance optimization rests on each sensor device and on the communication network. Do today's networks support appropriate geo-location, signal strength reporting, and antenna optimization?

3) Open Development Platform: Broadly-usable smart city infrastructure alone is not valuable unless paired with a rapidly-evolving software ecosystem. Recall that it was the million (or so) app developers who turned the phones-as-phones world into the phones-as-mobile-computers world. Smart cities should have, as a primary objective, the desire to attract these same million developers. Stale, purpose-built embedded programs for specific vertical applications in the smart city are not nearly so appealing as the premise of enabling motivated software engineers to develop city-scale apps that deliver value beyond the initially-imagined purposes for the smart infrastructure. In fact, we would do well to think of the city as a *platform* on which future apps will be built. And therein are several major problems.

First, we can imagine that such an open ecosystem could and should lead to a sharing economy for sensed data. This raises questions of how one might create a marketplace for information, how value-for-data might be formulated, and the extent to which this creates privacy issues. Second, an open ecosystem for programming raises the specter of the city's smart infrastructure being used improperly or, worse, being used against the city and its residents. As this new network becomes integral to daily life, concerns of resilience and the maintaining of network integrity become first-order considerations. Third, city apps are inherently cyber-physical where mobile phone apps, for the most part, are not. Smart city apps gather data from the real, physical world, process it, and then signal or trigger actions again in the real world. In so doing, signals cross the so-called Time Line (Figure 1) twice. The Time Line is the separation between the real world, where time has physical consequences and the cyber world, where time is simply meta-data. Programming and software engineering as taught and understood today rarely reflect a sufficiently deep understanding of the implications of Time-Line-crossings.

Finally, setting aside concerns of information privacy, city-platform abuse and time-programming complexity, the notion of multiplexing sensing elements in the city across multiple, separately-developed apps raises resource management challenges. In essence, the city will become a large computing aggregate, and questions of how to fairly share its resources will arise just as in the timesharing days of old.

Consequences: Innovation in the smart city relies on third party developers. They in turn will be attracted by an open platform. Enabling this requires solving fundamental problems of fairness, network integrity, value exchange, and time-aware programming [8]. Do the old rules and approaches for third party app developers apply?

4) Edgelessness: The power of the smart city is in its cyber-physicality. Timescales involved in detecting and avoiding motorist accidents are measured in milliseconds, and the sensing, computing, and actuating that take place in our smart cities must meet these expectations. The scales, costs, spans-of-control, and information sharing expectations call for synchronization mechanisms and latency management techniques for our traffic sensing systems that are a bit beyond the state-of-the-art, particularly with respect to how these are reflected in programs. Latency management by taking advantage of placing computation at the edge of the network is not a new idea–Compaq sold a line of commercial "edge of the network" servers in 2001.[3] Cisco has updated this concept and positioned it as bringing the cloud closer to the ground–so-called fog computing [9]. Satya and colleagues took a complementary view of pushing computation from (at the time) compute-impoverished mobile devices into processors that had

[2]http://pittsburghpa.gov/dpw/street-resurfacing
[3]http://www.serverwatch.com/news/article.php/1400281

cloud-like capabilities but that were proximate to the edge [10]–called cloudlets. All of these approaches focus on enabling the placement of computing near the network's edge.

We believe the bigger problem is enabling programmers to write single programs that can be automatically distributed and migrated in and between the devices, the network and the cloud–but without having to explicitly manage all of this partitioning complexity. In essence, the notion of an "edge" places a substantial burden on the shoulders of the programmer to decide how to cut his or her program and how to map it–both today and tomorrow when the relative computing capabilities of the devices, network nodes and cloud change.

We must embrace the realization that, in order to bring about a revolution in IoT and smart city computing for applications like traffic calming and others, we must reduce rather than increase the complexity of the programmer's task. We want the same million programmers who made mobile computing what it is today to adopt the smart city as their new platform. Today, they face a steep learning curve to write code that somehow coordinates and harnesses device, network, and cloud resources–needing to navigate the various edges without getting cut.

Consequences: We put forward the concept of edgeless computing, arguing that the large-scale economically-transformative change of smart cities in particular and the Internet of Things in general will only come when a programmer can write one program that harnesses device, network (edge), and cloud as easily as she can write a mobile app today, erasing forever the presumption that we must consider the network and device to be on opposite sides of an obtrusive interface (the so-called edge). Can today's programming environments and languages support this?

B. Today's Networks

An immediate and important question is *Do today's cellular networks meet these design criteria, or if not, could they be readily adapted?* Certainly, today's networks provide generally excellent coverage–a result of decades of careful network design and optimization. But this was done with the assumption that the terminal devices would largely be phones, held at human height, used for voice calling and data, and recharged every night. Devices in the smart city are different. Many, if not most, will be mounted on buildings, structures, signs, signals or (worse) the pavement. They will be used for simple telemetry applications (sending a few bytes of data when something changes). And, importantly, they will need to live a city-infrastructure lifetime (significant fraction of a decade), unwired, without being recharged. Any resemblance between these requirements and those of a cellphone is purely coincidental.

Cellular networks impose an assumption of network-centricity on devices. The requirement, even for the newest 3GPP protocols, for the device to stay connected and report in periodically thwarts efforts of IoT device designers to create truly low-power solutions. LTE MTC [11] and subsequent standards including NB-IoT [12] seem alluring to IoT architects looking for pervasive coverage. But the energy tax to simply stay connected to the network is still too high for devices that must last for five to ten years when operated from a coin-cell battery. And in many cellular markets, the premise of hundreds of thousands of cellular device subscriptions just for sensors is simply cost-prohibitive. Laying wires to solve the power problem also runs afoul of the city-friendly integration principle.

Looking at today's cellular networks, we must also ask ourselves if openness and edgeless computing can be brought about. It is all-too-apparent that today's cellular networks were not designed to be open development platforms. As an evolution of the venerable public switched telephone network, they were designed first and foremost to be hardened, reliable voice networks (it was not until 2009 that more data than voice transited commercial cellular networks [13]). Partially of necessity and partially of habit, cellular networks evolve slowly. Measured evolution was predicated on the costs and risks associated with vesting network and radio logic in hardware (an assumption that is rapidly becoming invalid in the face of software-defined radio). Nevertheless, current standards-setting and operational practices of cellular networks define a pace that makes rapid adoption of concepts like open networks and edgeless computing unlikely in the short term. The concept of embedding third-party programming in a carrier's network has, historically, been anathema to telecommunications network design–raising concerns about denial-of-service attacks, privacy, and security. While there may be approaches to bound and mitigate such risks (and, in fact, commercial IP-based data centers and networks do so every day), the cultural shift from yesterday's telecom central offices to internet practices may be next to impossible for legacy telecommunications providers.

We conclude that while cellular networks are *the* pervasive network, their shortcomings in terms of being closed, having inadequate support for ultra-low-power city friendly devices, and offering weak to nonexistent support for edgeless computing compel us to consider alternatives so as to accelerate innovation and hasten the arrival of the smart city.

C. Low-Power Wide-Area Networks

Low-Power Wide-Area Networks (LP-WAN) [14] are emerging as a new class of networks that are well-suited to the design principles for the smart city. These networks are built on novel, narrow-band communications technologies such as Semtech's LoRa chirp spread spectrum technology [4], Ingenu's random-phase multiple access technology[4], or SigFox's narrowband binary phase shift keying technology[5]. In all cases, the radio access network is optimized for low data rate transfers (kilobits per second) at very low duty cycles. These networks are being deployed in unlicensed spectrum. For example, in the USA, LoRa uses the 902-928 MHz Industrial, Scientific and Medical (ISM) band.

We examine LP-WANs against cellular networks in terms of device considerations, network considerations, and performance considerations.

1) Device Considerations for LP-WAN: For the smart city, LP-WANs offer significant advantages, at least at the technology level. Unlike cellular networks in which device power is determined by network-side timing and protocol considerations, thereby establishing a lower-bound on power consumption, LP-WANs are device-centric, leading to significant device power advantages over cellular. The typical modality for an LP-WAN device is to spend most of its lifetime asleep, waking on a trigger indicating the availability of new data, and only transmitting when the device has useful information to convey. For a device that must live for half a decade or more on, essentially, a standard, charged-once, cellphone battery, the ability to spend most of its time asleep is the only practical way to survive. LP-WANs are well-matched to this need.

The radio modems for LP-WANs are low in cost relative to cellular modules, operate in unlicensed spectrum, and need not pass through the lengthy process of testing for compatibility with a particular cellular operator's network. Because LP-WAN-enabled devices can be energy efficient, they can be small and totally wireless, making them in principle city-friendly.

Smart city applications, particularly for LP-WAN devices, raise important concerns for antenna performance. Because of its fundamental relationship to network design and device power, we give antenna performance special consideration in the context of a real application in Section II-E.

2) Network Considerations for LP-WAN: Outside of certain countries in Europe, LP-WANs are only sparsely deployed at present. As such, they don't compare favorably to cellular networks on the basis of pervasive coverage today. But to their credit, setting up an LP-WAN network is relatively more straightforward than setting up a cellular network. The network nodes themselves are small (cigar-box sized) and modestly priced (under USD 2000 each). The relatively low cost per gateway (compared to an equivalent-coverage eNodeB in an LTE network) opens the door to a middle-out network deployment that can be done incrementally.

While overlapping coverage of cellular sectors must be carefully engineered and controlled, in some LP-WANs such as LoRa, overlapping coverage is actually an advantage. With the ability for multiple gateways to hear transmissions from low-power devices, a measure of redundancy is introduced. With some care, localization of devices is possible when three or more gateways receive the same LoRa packet, offering a coarse-grained alternative to on-device GPS.

Moreover, because the backhaul bandwidth from an LP-WAN gateway is limited by the low-bandwidth radio network itself, these gateways can be connected to the internet via modest-speed cellular connections, in effect making LP-WANs an overlay network on top of cellular. These factors enable rapid establishment of LP-WANs for smart city (and other) applications. An overlay approach allows LP-WAN networks to be built out incrementally, and the cost of the cellular connection can be amortized over thousands of LP-WAN nodes.

The fact that LP-WANs are not established actually has a further advantage. With no standing assumptions about network architecture, the creation of an open, programmable network is a real possibility. Integrating commodity computers with LP-WAN gateways and open source virtualization tools creates an exciting possibility of a new kind of network that is very well matched to the smart city design criteria.

3) LP-WAN Summary: In *The Innovator's Dilemma*, Christensen [15] describes disruptive technologies as

> ...technologies that result in *worse* product performance, at least in the near-term... Disruptive technologies bring to a market a very different value proposition than had been available previously. Generally, disruptive technologies underperform established products in mainstream markets. But they have other features that a few fringe (and generally new) customers value...

[4]http://www.onrampwireless.com
[5]http://www.sigfox.com

We see LP-WANs as a potential disruptor to today's mainstream cellular networks. They underperform in mainstream markets, but they offer low-power operation, inexpensive buildout, and some hope for edgeless computing–especially appealing to the emerging market of the smart city (and, more generally, to the IoT). Not all disruptors succeed. But those that do often completely overturn existing markets and existing players. We believe LP-WANs are particularly well-suited for smart cities when they are deployed openly and augmented with core services and edgeless computing capabilities.

D. Crowdsourcing Information in the Smart City

In an earlier project aimed at reducing the cost of gathering data for early earthquake warning using ordinary mobile phones [16], we were awakened to the power of the crowd. Motivated by this, and considering the (disruptive) potential of LP-WANs, we are struck with an interesting observation. The traditional smart city approach is rather top-down, with a presumption that some number of specific vertical applications will be proposed, specified, funded and built out, and possibly at high (prohibitive?) cost due to networking complexities, the need to provide power and/or networking to sensors, and all the issues related to permitting and financing such work.

What if we dispensed with the concept of a bespoke built infrastructure for smart city functionality and, instead, incrementally built an LP-WAN, augmented with computing embedded in the network along with some basic cloud-side services for device enrollment? This smart city ad-hoc platform could be opened to non-specialist application developers with the challenge to build novel smart city apps. While at first blush this may seem outlandish, the possibilities are quite real.

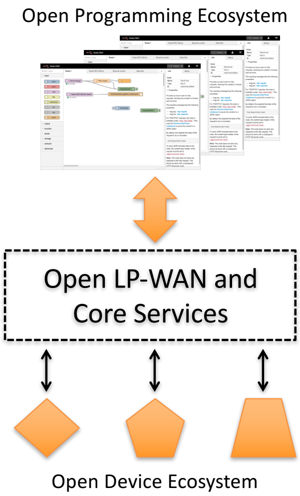

Fig. 2: An open LP-WAN with core services and edgeless computing is the essential "middle" of the crowdsourced smart city.

As has been seen in the recent movements toward single-board-computer-enabled projects (*e.g.,* Arduino[6], Raspberry Pi[7]), a common hardware platform and suitable example applications enable significant innovation. We see it as beneficial to consider how smart city infrastructure should be made open to hardware and software developers and to encourage wide participation in an incremental build-out of enablers for novel applications.

We call this broader vision *crowdsourced smart cities* because it seeks to build on the power of the crowd to instrument cities and to develop the clever software applications that will have practical value to city residents. We take as our first realistic application the problem of *crowdsourced traffic calming*.

This kind of middle-out concept–creating a common network that will adapt to many types of hardware sensors and that enables many applications–is depicted in Figure 2. By seeding a city with a basic LP-WAN and complementing the raw network with core services (Figure 3) for basic data storage, retrieval and visualization, we enable both the simple attachment of a wide range of new sensor types to the city, particularly inexpensive, wireless traffic sensors, and a corresponding open programming ecosystem to extract insights and value out of the traffic data so collected.

Fig. 3: Minimal LP-WAN services to support crowdsourced smart cities.

[6]https://www.arduino.cc/
[7]https://www.raspberrypi.org/

The common middle is a foundation for sensed-data storage, processing and visualization–elements that would otherwise have to be re-created for each smart city project. By sharing these, we accelerate development. By providing the means for participants in the crowdsourced smart city, especially non-specialist participants, to make data available to one another, we potentially open new kinds of applications that no one might have been able to undertake.

The possibilities of such large-scale information sharing must be balanced against privacy and security concerns and the need to maintain resilience of the network itself.

E. Bringing the elements together: Crowdsourced Traffic Calming

Both traffic calming and the broader notion of modeling and understanding traffic flows in the smart city can benefit significantly from a systematic means for continually collecting and analyzing traffic data in real time. This necessitates a pervasive network of sensors, time-stamped geo-referenced traffic readings, a network for collecting these, and the means to logically aggregate data and interpret the inputs as traffic flows over time. Our approach is to create a system made up of (a) a small sensor board that includes an LP-WAN radio, transducers, and other devices that can be embedded in city infrastructure; (b) a simple LP-WAN network to provide these sensors with connectivity to the internet; and (c) a web-based data collection, processing and visualization toolkit supporting both traffic analysis and sensor management.

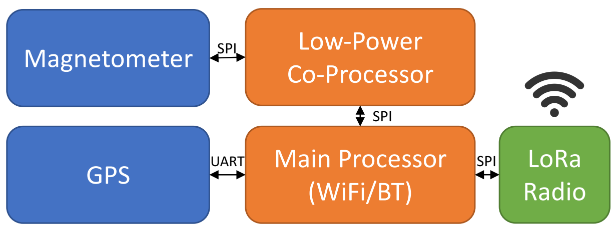

Fig. 4: Main components of the Traffic Sensor

Our sensor board includes a three-axis magnetometer, a small processor, an electronic serial number (ESN), a battery and a radio subsystem including Bluetooth Low Energy, WiFi, and a LoRa LP-WAN. The board has been engineered for a five-year lifetime with a single battery. While the board includes solar recharging capability, we don't count on harvested energy to achieve the lifetime target. Previous studies have shown the effectiveness of using magnetometers for sensing traffic [17]–[19].

The board can be built into various roadway fixtures. But we are focusing specifically on packaging this sensor board into a common roadway lane marker because some cities (and suburban areas, and rural areas) already have practices for installing these. Having the sensor in close proximity to the vehicles themselves makes vehicle detection relatively straightforward.

Our starting point was the so-called Botts Dot,[8] invented by Dr. Elbert Dysart Botts who, as an engineer with the California Department of Transportation (Caltrans), sought to reduce accidents by making lane lines more visible, particularly in the rain. Botts Dots have evolved, and the more popular form is the Stimsonite-type[9] roadway marker. Generically, such devices are referred to as Raised Pavement Markers (RPM). These are often seen with colored retro-reflective tape or insets.

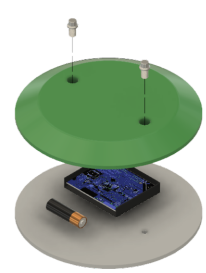

Fig. 5: CMU's Traffic Sensor – an intelligent RPM for counting vehicles and measuring speed.

1) The Sensor: We've taken the concept of a passive RPM and added our board to it, creating a smart Traffic Sensor (Figure 5). RPMs are already pervasive in many cities; we are designing our Traffic Sensors to be mechanically interchangeable with existing RPMs. RPM packaging imposes stringent constraints on

[8]https://en.wikipedia.org/wiki/Botts'_dots
[9]http://www.ennisflintamericas.com/downloads/dl/file/id/38/product/1038/brochure_model_101_rpms.pdf

size/shape, mechanical loading, water-tightness, inaccessibility post-installation (we call this the OHIO principle–we can Only Handle It Once), thermal stresses (-20°F to 150°F or worse), and the occasional snowplow. Size and shape constrain the dimensions of the all-important LP-WAN antenna which we consider below.

To effectively monitor traffic flows, our Traffic Sensors can be placed judiciously, for example, at points leading into and emerging from intersections. Where speed on a long roadway is a concern, a Traffic Sensor can be placed exactly where the measurement would be most meaningful.

In order to correctly geo-reference its readings, the location of each Traffic Sensor must be recorded. This can be accomplished (a) at installation time by manually recording each Traffic Sensor's ESN, its latitude and its longitude, (b) post-installation using a drive-by technique with a smartphone app to capture wirelessly-beaconed ESNs and to records the phone's corresponding GPS position, (c) by the LP-WAN network or (d) by the Traffic Sensor itself (we've built in a GPS receiver for this purpose). We are exploring the accuracy by which the network can localize Traffic Sensors. Network-based localization, if adequately accurate, could make geo-referencing of Traffic Sensors transparent to city practices and could obviate inclusion of a GPS module.

The OHIO principle reminds us that developing and programming a Traffic Sensor is not unlike developing and programming a mission to Mars–once we launch the Traffic Sensor (epoxy it, or 500,000 of them, to the road surface), there is no going back. This suggests, among other things, that our Traffic Sensor and similar smart city sensors need to be re-programmable over-the-air. We are designing our Traffic Sensors to accept parameter changes and/or incremental software updates on a scheduled, broadcast basis over wireless networks.

2) Antenna Design and Self-Optimization: The economics of wireless networks, generally, rely on low cost per area covered. Each fixed gateway should cover the largest possible area. The physics that drive coverage involve topography, structures, the way the information is coded, the noise in the radio channel, and the antenna subsystem design. While novel modulation techniques for LP-WANs provide valuable coding gain, the physical constraints imposed on LP-WAN devices in the smart city work against good signal propagation. Figure 6 captures the essential elements, summarized by Equation 1:

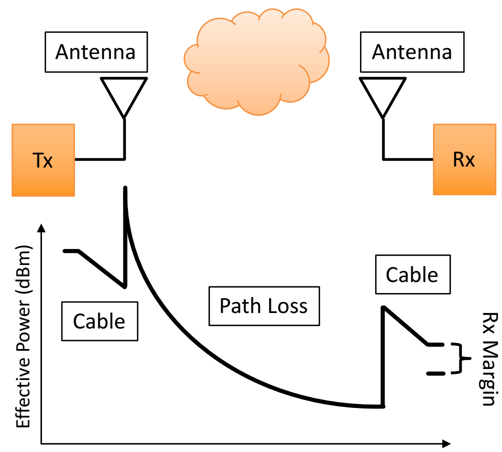

Fig. 6: Signal strength at the receiver is a function of transmitter signal strength, cable losses, antenna gains (or losses), and free space path loss. With power being capped by regulation and receive sensitivity set by technology, maximizing distance is done by optimizing the antenna subsystems.

$$P_{rx} = P_{tx} - P_{txcbl} + P_{txant} - PL + P_{rxant} - P_{rxcbl} \tag{1}$$

where P_{tx} is the output power of the transmitter (capped by regulation), P_{txcbl} and P_{rxcbl} are losses attributable to the cables at the transmitter and receiver, P_{txant} and P_{rxant} are the gains (or losses) of the transmit and receive antennas, and PL is the path loss between the antennas. P_{rx} is the resulting power available at the input to the receiver and must be above the receiver's sensitivity (governed by coding design and hardware considerations).

The distance that can be covered, then, is only that which, given the antenna, transmitter, receiver, and cable characteristics, keeps the signal at the receiver above its minimum level. This can be approximated by the idealized free space path loss term, a function of frequency (f) and distance (d), and is given (in decibels) by Equation 2:

$$FSPL_{dB} = 20log_{10}(\frac{4\pi df}{c}) \tag{2}$$

Path loss, antenna design, and the relationship to distances and areas an LP-WAN gateway can cover are addressed in our companion paper [20].

City friendliness readily translates into "no external antenna" for many applications. Moreover, device placement (on infrastructure, on buildings, on pavement surfaces) compels antennas to be in ground proximity, leading to signal energy aimed sub-optimally rather than at the nearest gateway. We thus face significant antenna-related challenges:

Packaging: LoRa's frequency of operation is a first-order consideration in designing a suitable antenna. A half-wave wire dipole antenna in air at 915 MHz would be approximately 6.1" across. This will not fit inside a typical RPM.

Ground proximity: A second problem is the effect the pavement itself has on the antenna's pattern. Mounted in an RPM, a horizontally-polarized antenna's height above ground will only be a small fraction of a wavelength–leading to a pathological "straight-up" radiation pattern and increasing the so-called elevation angle (the angle from horizontal at which the antenna's pattern is at a maximum–see Figure 8). Such an antenna is euphemistically referred to as a *cloud burner* because the bulk of the energy is simply dissipated as heat in the atmosphere instead of yielding a strong signal at the receiver. Can better antennas be developed to mitigate this problem?

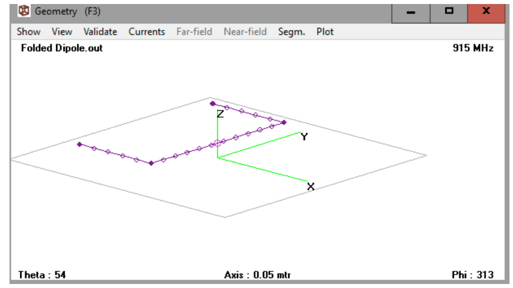

Fig. 7: A 915 MHz horizontally-polarized dipole antenna, folded to fit inside an RPM.

Aim: The third issue of concern is how the RPM is aimed. If the antenna's azimuthal pattern (horizontal plane) is nonuniform, then the orientation of the Traffic Sensor relative to the gateway may have a detrimental effect on received signal strength.

What are the practical impacts of these issues? Figure 7 shows the geometry of a horizontally-polarized 915 MHz antenna, folded to fit an RPM. Its performance, both as a result of folding and as a result of ground proximity, is shown in Figure 8. The compromised geometry, at a 5° elevation angle, accounts for >10dB of effective signal strength loss.

Less-than-ideal antenna performance–due to geometry, ground proximity and/or aim–will lead to higher energy-expended-per-bit-transmitted and shortened battery life. The premise of periodically replacing half a million batteries in a smart city's traffic infrastructure serves as a motivator for improving RF performance in other ways.

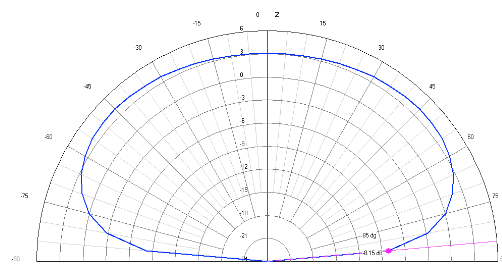

Fig. 8: Geometries and placement of smart city devices can compromise antenna design. In this case, a folded horizontal dipole antenna packed into a small sensor mounted on pavement would exhibit near-vertical-incidence behavior (main lobe points upward). Gain at a low elevation angle (in red)–such as would be the case of the device transmitting toward the gateway–is 10-20 db below the main lobe. We are developing alternative antenna strategies to mitigate this effect. [20]

Because we can't count on installation-time optimization, the burden must fall to the devices and the network to be self-optimizing in terms of RF performance. Likewise, installation is simplified and cost is reduced by tasking the devices and the network with accurately recording the position of each sensor post-installation.

We are able to at least partially address these issues through antenna design that optimizes the low elevation angle demanded of our Traffic Sensors and provides an adaptive means for the Traffic Sensor to beam-steer its signal, using a combination of gateway signal strength measurement and a beam-forming antenna array. Our deeper treatment of our approach to antenna optimization is given in [20].

Energy conservation can also be enhanced by adapting power levels, information encoding, and frequency of transmission, subject to the constraints of the overall system's design objectives (such as timeliness and resolution of measurements). We explore these issues and their relationship to Traffic Sensor battery lifetime in [21].

3) System Management and Analytics: System management and analytics are necessary tools for controlling the behavior of devices and analyzing the data flowing through the system. Being able to do this rapidly and flexibly is preferred when running frequently changing experiments. The time to build up functionality to support experiments should be minimal compared to the time it takes to set up a custom, purpose-built solution.

Given these principles, we chose to rely on available open source tools that support the needed functionality and customizations. The first is Node-RED [22] - seen in 9a - provides a flexible, visual-based programming environment, which is useful for wiring up data flow diagrams and easily changing settings for physical nodes. Second is InfluxDB

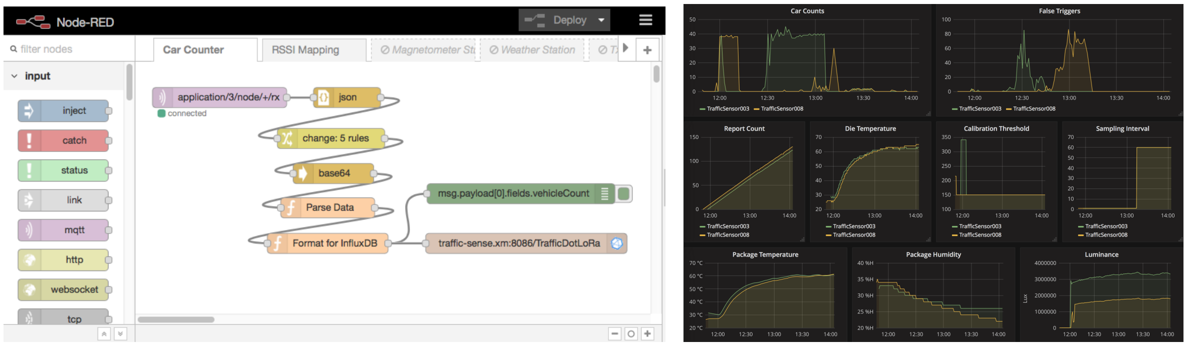

(a) (b)

Fig. 9: System management and analytics tools. (a) Node-RED is used for programming and prototyping data flows between devices and data ingestion tools. (b) Grafana is a flexible analytics dashboard for displaying time-series data.

[23], a time series database engine that makes it easy to align and correlate time series data. Last is Grafana [24], Figure 9b, which provides a rapid development of dashboards and visualizations for time series data. All of these tools contribute to the rapid management devices and real-time analysis of events.

III. Methodology and Findings

In this section, we present our methodology for evaluating our TC enablers and results from the salient aspects of the project: mainly vehicle sensing and characterization, evaluation of the LoRaWAN Network deployment and development of a Traffic visualization and Analytic system.

A. Vehicle Sensing and Characterization

The first challenge in creating a traffic measurement system is developing a sensor that will detect the presence of, and possibly extract other valuable information about, passing vehicles. MEMS magentometers have been used successfully in previous studies [17]–[19]. Our sensor extends this work and seeks to extract additional information from the magnetometer's signal through low-power, on-sensor processing. In the following sections, we review our experiments to characterize vehicular magnetic signatures captured by our Traffic Sensor.

1) Vehicle Magnetic Signature Data Set: To enable development of algorithms for sensing and characterizing vehicles, we designed an experiment to capture detailed magnetic signatures, ground-truth speed and vehicle type on an off-road test track. A vehicle signature is a tuple

$$SIG = (type, speed, \{mag_x(t), mag_y(t), mag_z(t)\}) \tag{3}$$

where $type$ is drawn from a set of four distinct sedans and four distinct SUVs, $speed$ is based on measurement, and $\{mag_x(t), mag_y(t), mag_z(t)\}$ is the set of samples from the magnetometer in x, y, and z, running at 400 Hz. Ground-truth $type$ is derived from images taken at the time the sample was collected. Similarly, ground-truth $speed$ was taken from a radar gun aimed directly at the oncoming vehicle. The experimental setup is detailed in Figure 10. Three of our Traffic Sensors with magnetometers were arranged in the center of the travel lane. In total, we collected 144 distinct signatures. Figure 11 shows magnetometer traces from three different vehicles.

2) Observations: We observe from Figure 11 that the presence of a moving vehicle is easily detected as compared to the stationary background field. Detecting and isolating these signals from the background is possible with thresholding techniques. Comparing different vehicles, we observe that each vehicle induces a unique signature in all three axes of the magnetic field. Furthermore, we see that at different speeds, these signatures are distinguishable but appear stretched or compressed as a function of speed. With these observations, we designed algorithms to identify a particular vehicle and estimate its speed based.

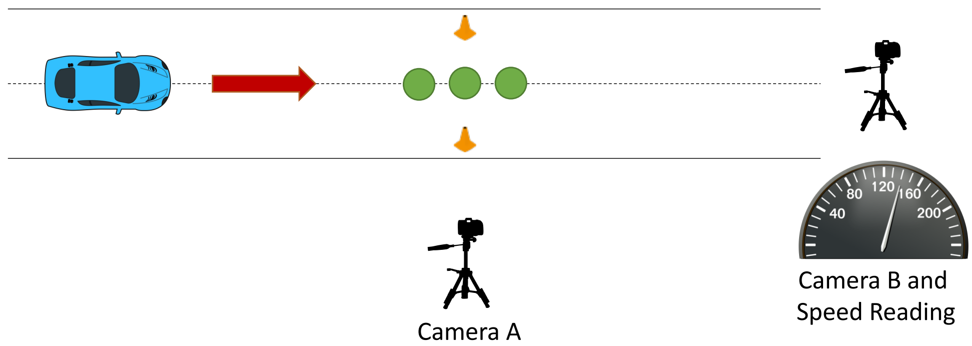

Fig. 10: Experimental Setup: Three magnetometer sensors were placed along the center of the travel lane with traffic cones marking the location of the sensors. The speed of the vehicle is captured by a radar gun as the vehicle passes over the sensors. Two video cameras capture these crossings from in-front and from the side of the crossing.

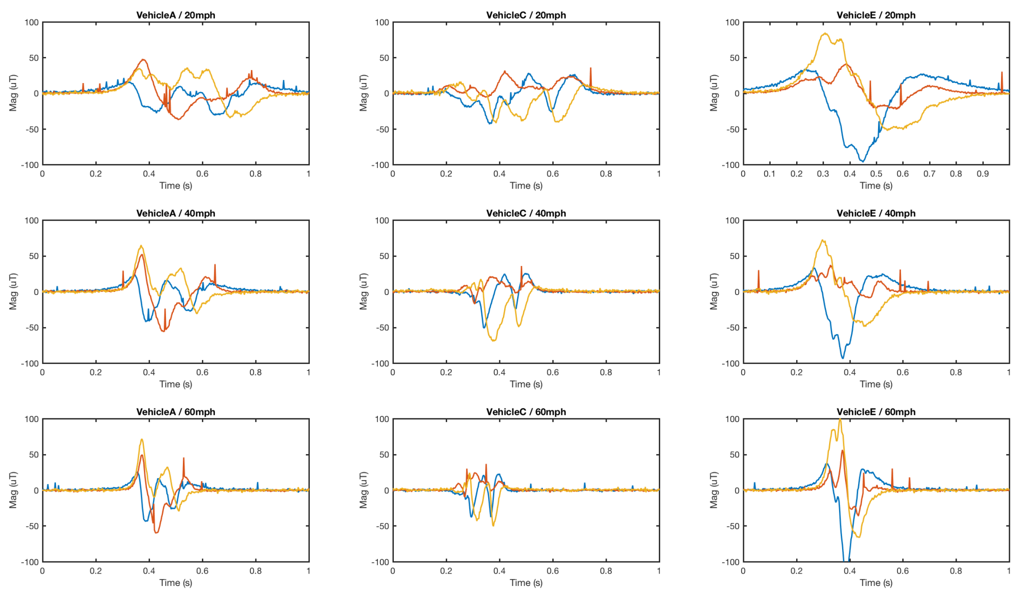

Fig. 11: A sample of 3-axis magnetometer waveforms from three different vehicles moving at three different speeds. Observing the signatures, we can see that each vehicle provides a signal that is distinct from the others. For a given vehicle, the signature is stretched or compressed in time as a function of speed.

3) Presence Detection Algorithm: Simple traffic flow studies are enabled by extracting vehicle presence from the magnetometer traces. We can pick out a moving vehicle from the static field readily. When first installed, and periodically, each sensor must measure and remember the non-disturbed, or background, magnetic field. Perturbations are then readily detected. Not all perturbations are vehicles, so additional processing is necessary to extract the characteristic changes in the three axes to separate signal (passing vehicle) from noise.

The algorithm begins by characterizing the static magnetic field. We measure the 3-axis means and standard deviations of the magnetic field with a sliding time window. We then determine a vehicle presence event when a successive window contains energy that is several standard deviations higher than that of the static noise floor. Consecutive high energy windows comprise a single detection event.

While this approach can easily detect any perturbation of the surrounding magnetic field, our algorithm must minimize random detection events and maximize the actual presence of vehicles. This can be performed by evaluating the length of the detected signal. By our observations, vehicles running at the highest reasonable speed (60mph) will result into at least 0.2s of magnetic perturbation, given the average length of vehicles.

By requiring at least 0.1s of magnetic perturbation, we can limit detection events that do not pertain to vehicular activity. Meanwhile, a detection event spanning several seconds could mean that a vehicle has stopped in the presence of the sensor–indicating stalled traffic. Extended static shifts trigger re-calibration of the magnetometer to the new background field.

4) Vehicle Identification and Speed Estimation: Vehicle identification and speed estimation are important to traffic calming to measure the utilization and safety of traffic corridors. Using our initial observations, we developed a base algorithm to identify a vehicle and estimate speed as described in Figure 13.

The algorithm utilizes a reference signal from a target vehicle moving at a known speed. Events that correlate highly with a scaled version of the reference signal signify a match, while the scale estimates the speed at which the vehicle passed by using Equation 4.

$$EstimatedSpeed = \frac{ReferenceSpeed}{Scale} \tag{4}$$

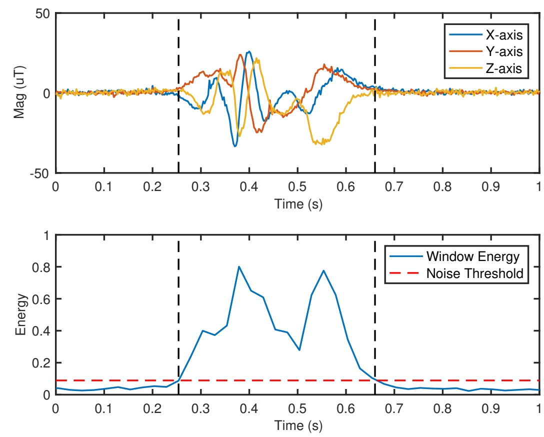

Fig. 12: Event detection: By measuring the signal energy in small windows of time, we can detect the presence of an object distorting the static magnetic field around the sensor. Setting a threshold higher than the noise floor allows us to isolate the signal and further process it.

Using fixed thresholds and reference events from each vehicle running at 30 mph, the algorithm is able to achieve results shown in Table I, where results are described in terms of precision, recall and F1 score (harmonic mean of precision and recall).

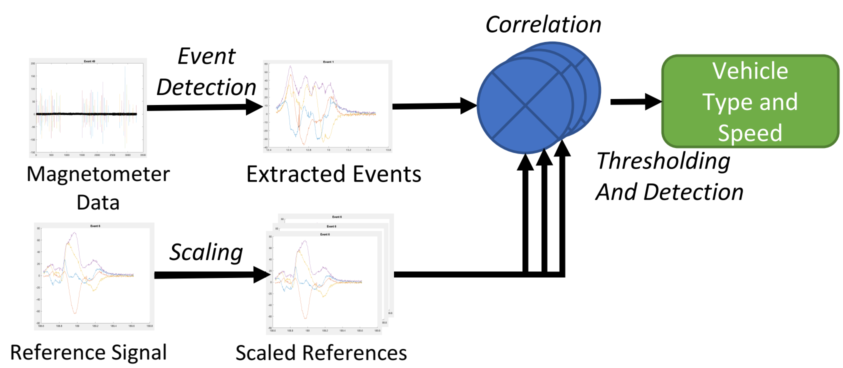

Fig. 13: Signal processing pipeline for determining vehicle type and speed from a reference signal.

These results show that the algorithm correctly identifies which exact vehicle passed over the sensor with an accuracy of 91%. This includes being able to differentiate between vehicles of the same body type, *e.g.,* sedan A vs. sedan B, while running at different speeds.

Furthermore, the algorithm correctly estimated speed to within 5 mph of the reported value for 87.5% of events running up to 60 mph. The algorithm performance degrades at speeds higher than 60 mph (attributable to the magnetometer's sampling rate–faster sampling would mitigate this issue). For measuring the effect of traffic calming interventions in, say, neighborhood streets, nominal speeds would likely be below this. It can also be seen in Figure 14, that the algorithm performs acceptably across all vehicle types, with the exception of Vehicle #8.

5) Future Work–Vehicle Sensing: While the approach provided works well in identifying a specific vehicle, it requires a reference waveform and speed from the target vehicle. Further characterization of the dataset can result in improved classification algorithms to determine vehicle type and accurately measure speed using one device without the need for a reference waveform. Some effort has gone into using machine learning algorithms to classify vehicle types. This study has motivated our follow-on work to collect larger and broader datasets and to explore

Vehicle	T.Pos.	F.Neg.	F.Pos.	Precision	Recall	F1 Score
1	17	1	0	1.000	0.944	0.971
2	17	1	2	0.895	0.944	0.919
3	17	1	0	1.000	0.944	0.971
4	17	1	9	0.654	0.944	0.773
5	15	3	2	0.882	0.833	0.857
6	17	1	0	1.000	0.944	0.971
7	13	5	0	1.000	0.722	0.839
8	18	0	0	1.000	1.000	1.000
Overall	**131**	**13**	**13**	**0.910**	**0.910**	**0.910**

TABLE I: F-1 score of correctly identifying a vehicle based on correlations with a reference waveform. The algorithm is able to correctly identify a vehicle 91% of the time versus all other vehicles in our dataset.

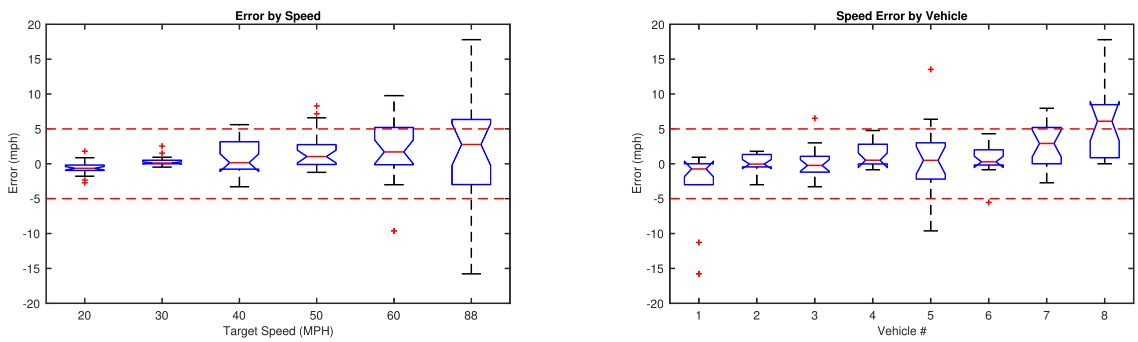

(a) (b)

Fig. 14: Speed estimation results classified by (a) target speed and (b) vehicle. Results show acceptable performance for speeds below 60mph to within 5mph of the target speed and comparable performance across vehicle types, with exception of Vehicle #8.

the use of machine learning techniques to associate magnetic signatures with vehicle types and speed classes. This is a matter for follow-on studies.

B. LoRaWAN Network Deployment and Coverage Mapping

LoRaWAN is a new type of wireless network for smart city applications. Various studies have reported LoRa communications over specific, and in some cases, idealized terrains of 10's of km. In realistic deployments, the ability to place devices anywhere with an expectation of reliable communications is an over-riding consideration to raw distance records. For a realistic deployment of a communications network to support traffic measurement studies, measurement techniques for network coverage are essential. To that end, we developed a suite of mobile- and web-based tools for collecting LoRa received signal strength at various locations to analyze the network coverage achieved with a single gateway. The following sections describe the tools and results gathered from the system.

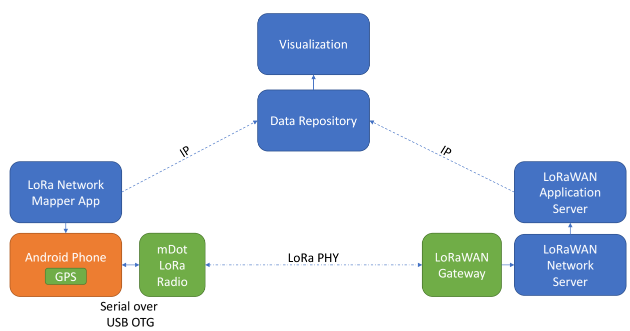

Fig. 15: Architecture of the LoRa Network Mapper application. An Android application periodically communicates and samples LoRa signal strength from an associated gateway. Geolocation is provided by the phone's internal GPS and is sent along with each sample. Signal strength from both gateway and end-device is submitted to a data repository where it is processed and visualized.

1) LoRa Network Mapper Application: The LoRa Network Mapper is a suite of tools for measuring and analyzing the coverage of a single LoRaWAN gateway.

This consists of an Android-based application to drive a LoRa radio, a LoRaWAN network deployment and a server-based processing pipeline for analytics and visualization, as described in Figure 15. The primary measurement tool is a LoRa radio connected to an Android phone. We designed a custom Android app to periodically send LoRa packets and measure the received signal strength of the gateway's responses along with the GPS location of the phone. The app also provides real-time information on collected data to guide the user to areas that require sampling. All measurements, including packets received on the gateway, are recorded and logged in a data repository. Coverage over a wide area is then interpolated using the so-called natural-neighbor algorithm[10].

2) Palo Alto LoRa Gateway Coverage: A city poses a challenging terrain for wireless signals. It is important to measure actual wireless coverage conditions as a figure-of-merit for gateway placement. Working with the City of Palo Alto, we deployed a LoRa gateway with an omnidirectional antenna on the roof of Palo Alto City Hall. Using our network mapper application, we gathered signal strength measurements in a wide area around city hall and visualized it. This is shown in Figure 16.

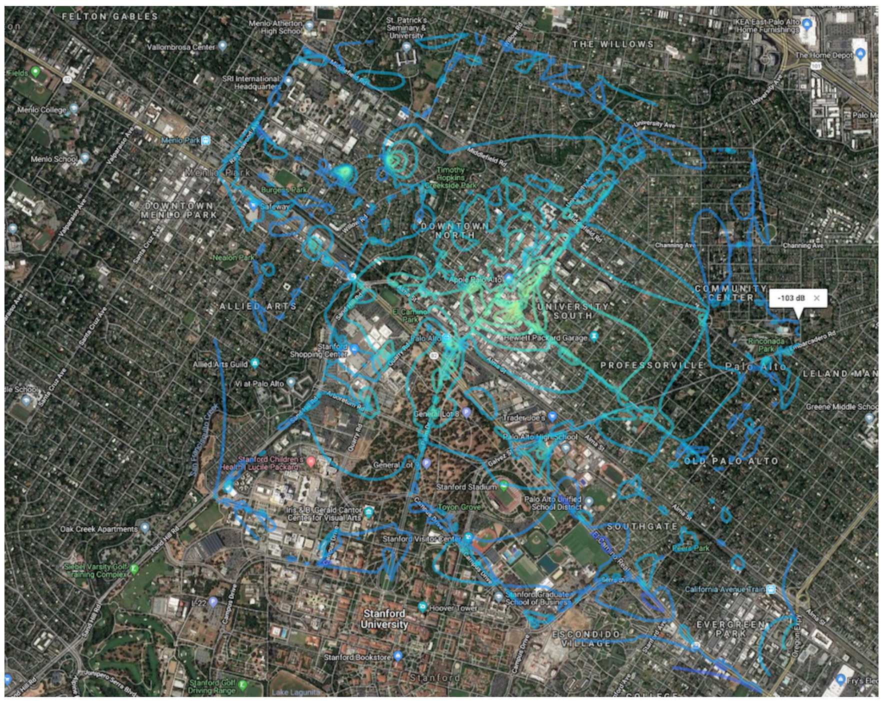

Fig. 16: LoRa received signal strength coverage map of downtown Palo Alto showing interpolated contour regions of signal strength in 5 dB steps (yellow is strongest and blue is weakest). At a radius of one mile from the gateway, the signal strength is measured to be at a quite acceptable -103 dBm. (Map Source: Google Maps)

Our results show that a single gateway is more than enough to cover a radius of a mile from the gateway in this sort of urban terrain with margins to spare. Such an area can include a few hundred road segments. If we place a sensor in each travel lane of each segment and conservatively assume two travel lanes in all road segments, a single LoRaWAN gateway can provide good RF coverage for up to 1,000 sensors.

While LoRaWAN gateways claim to serve up to 10,000 devices, the density and duty cycle requirements of sensor applications will ultimately limit the number of devices served by a single gateway. But with the relatively low cost of gateways, adding another one in an area has the benefit of increasing coverage, capacity and resilience of the network. Such networks scale up gracefully in a city- and cost-friendly manner.

3) Future Work–Network Deployment: The network mapper application has proved useful for various LoRaWAN deployments and has been made available to Comcast as a research tool for the characterization of their machineQ LoRa network. We foresee pportunities to expand the capabilities of the app, such as including support for multiple gateways, real-time coverage and traffic analytics to enable deeper studies in the long-term use and deployment of LoRaWAN-based systems.

[10]https://en.wikipedia.org/wiki/Natural_neighbor_interpolation

C. Communication Energy and Antenna Optimization

A wireless data collection system is no stronger than its weakest link. Shape, size, and materials for unobtrusive, robust sensor devices in the environment, particularly on road surfaces, impose substantial constraints on the antenna system which, without due attention, becomes the weak link. We devoted considerable attention to the fundamental problems of antenna design for roadway-mounted sensors. Poor antenna design translates directly into either increased system cost (more wireless gateways per square mile), reduced sensor lifetime (to overcome a poor antenna, devices have to put more power into transmitting, draining precious battery resources), reduced reliability, or some combination of these factors. In this section, we discuss the limits our device face and the steps we've taken to mitigate and improve upon these conditions.

1) Physical Constraints and Ground Proximity: As we design our devices to be drop-in replacements for raised pavement markers, we limit ourselves to specifications that are deemed safe for all passing vehicles, bicycles and pedestrians-alike. The overall device, including packaging, should be constrained in height, slope gradient and in all respects should match the physical specifications of existing RPMs.

We have started with the presumption that core drilling to install roadway sensors incurs substantial cost and is a dis-incentive to broad deployment. Instead, we elected to design our sensor to be surface-mounted just like traditional RPMs. We've also chosen to limit the height of our devices to 0.75". Load-bearing and waterproofing considerations consume roughly one third of that limit for structural integrity under load[11]. This restricts the space within for electronics, batteries and antennas (*e.g.,* a single AA-sized battery is typically 0.55" in diameter). Working collaboratively with the City of Palo Alto, we also took on the constraint of maintaining a maximum up-slope around the periphery of the device to 30° to minimize flip-over potential of passing bicyclists. Figure 5 is a rendering of the resulting package along with the components expected to be housed inside.

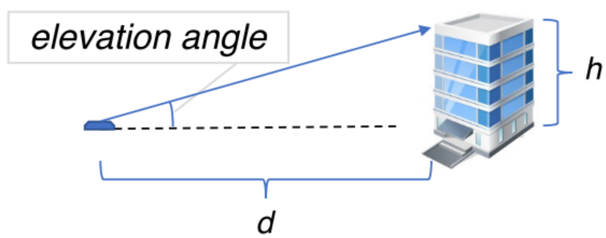

Fig. 17: Elevation angle between a traffic sensor at ground level towards a wireless gateway placed high atop a building. At long distances, the elevation angle can be quite low requiring efficient RF propagation near the horizontal plane. [20]

Having our devices at ground level compromises the performance of the antenna, sending most of the energy straight up, instead of towards the network gateways. Elevation angles in typical configurations are low as shown in Figure 17. Using antennas that most obviously fit the package, such as the horizontally-polarized dipole antenna shown, will result in signal loss as seen in Figure 8.

2) Packaging and Antenna Co-design: To optimize antenna performance with the physical constraints and ground proximity, we sought out to design the antenna in conjunction with the packaging. In our study, by co-designing the packaging mechanical structure to contain reflective elements, we can improve the gain of our device antennas by up to 6 dB with respect to a horizontal dipole [25].

Another promising approach is to use vertically polarized antennas. However, it is challenging to achieve efficient vertical antenna performance due to height constraints (at 915 MHz, the device's height is approximately 6% of a wavelength–far short of the desirable 1/4 wavelength for a traditional vertical dipole). Kong *et. al.* [26] created a vertically polarized spherical meandering antenna that has excellent antenna gain at low elevation angles in simulation, as seen in Figure 18a. We have taken this study and applied it to LoRa Frequency Bands (902-928MHz), achieving comparable simulation results as shown in Figure 18b. It remains as future work to solve the coupled RF and mechanical system design problem to create an enclosure that houses this antenna and simultaneously satisfies the other packaging constraints.

[11] Our sensor is designed to withstand repeated impact and down-force loads up to 6000 lbs.

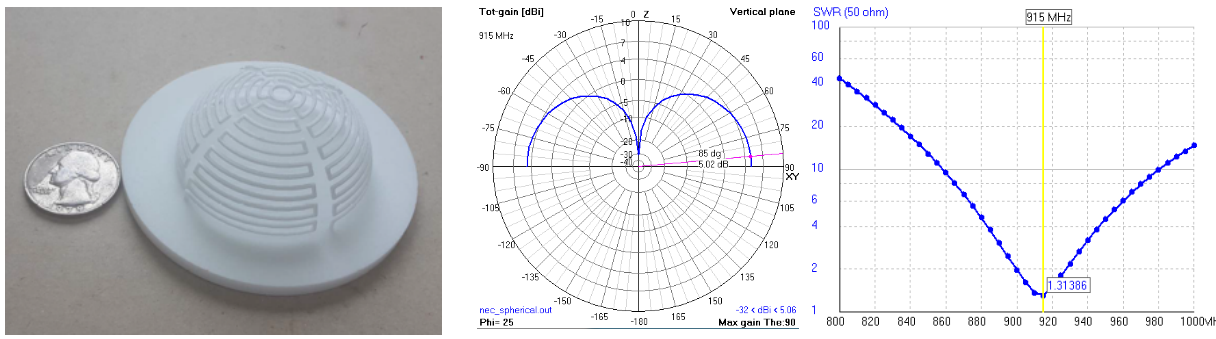

(a) (b)

Fig. 18: Kong's low height, vertically-polarized antenna in ground proximity with good performance. (a) the original antenna tuned for 720 MHz [26]; (b) our adapted version of the antenna tuned for 915 MHz, showing a gain of 5.02 dB at 5 degree takeoff angle.

3) Future Work–Antenna Optimization: We have studied the limitations and constraints of wireless traffic sensors deployed on pavement. Our results show that co-designing a vertically-polarized antenna with the packaging's mechanical structure will yield improvements in communication efficiency over more traditional antenna approaches, yielding important energy savings for the sensor. Further work is indicated to carry these simulations forward to an enclosure design that meets all RF and mechanical requirements.

D. Traffic Visualization and Analytics

Visualization and analytics are the tools necessary for interpreting the conditions of traffic flow in a given area. Enabling fast and real-time access to data makes it easier to make decisions on how to best proceed forward with any traffic calming intervention.

In the following sections, we detail the proposed requirements and need for visualization, discuss our current system iteration and propose future iterations of the system.

1) Requirements: Through discussions with Traffic Management Officials from the City of Palo Alto, we defined the metrics seen in Table II as the points of measurement that need to be analyzed. These values can be measured by counting the number of cars that pass through a segment, measuring their speed and classifying them into a vehicle type.

Name	**Definition**
Throughput	Number of vehicles passing through a road segment at any given time.
Utilization	Percentage of capacity at which the road segment operates.
Travel Speed	Distribution of vehicle speed along a road segment.
Vehicle Class	Distribution of vehicle types along a road segment.
Traffic Flow	Percentage of traffic that flows from one segment into another.

TABLE II: Metrics for Traffic Flow Analysis

The data query interface must be presented on a map to easily select and correlate the performance of one segment compared to another. Furthermore, it must be possible to view historical data on a segment as well as compare different time ranges to determine performance changes from one road configuration to the next.

2) Map Visualization and Simulation: With the requirements set above, we have built several tools for visualizing the defined metrics and displaying them on top of Google Maps. The data used is generated from SUMO [27], an urban mobility simulator in which we can define virtual sensors to act in place of our traffic sensors.

Figure 19 shows the interface displaying a selected road segment and comparing throughput across two time ranges. Data available is separated into two different travel directions, and if available, separated into per lane level

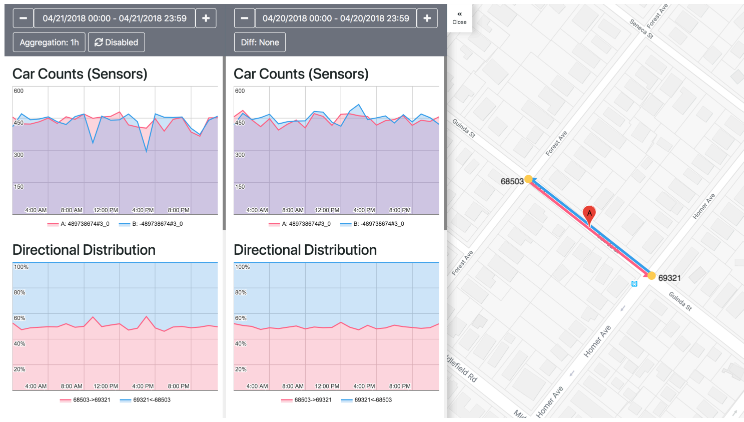

Fig. 19: Road segment visualization comparing throughput across two time-ranges.

information. Additionally, if road segment capacity is available, we can display the utilization of the road segment apart from throughput.

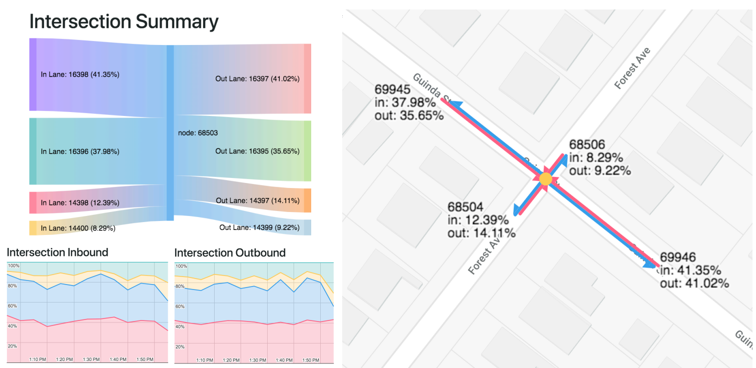

Fig. 20: Intersection visualization for Traffic Flow analysis.

Figure 20 shows the interface displaying a selected intersection, on which we display Traffic Flow information. The Sankey Diagram[12] (left) provides information on the distribution of number of vehicles entering and exiting the intersection from adjoining road segments. Given the available data, we can only display the percentages of vehicles entering and exiting the intersection based on the total contribution of all adjoining road segments. It will take further analysis to determine how much of a road segment's input is attributable to a previous segment's output.

3) Future Work–Visualization and Analytics: Our exploration of map visualization frameworks uncovered that current available tools do not provide adequate information necessary to define lane level data. Additionally, city management officials have specific definitions of road segments that are not captured in available tools. It may be worth exploring these needs.

Further work must also be done to fully develop the visualization system to support all use cases sought by the city, as well as integrate sensor data from deployed traffic sensors.

IV. Outcomes and Conclusion

We have created a Crowdsourced Traffic Calming System that enables non-specialists to take part in systematic measurement and evaluation of traffic calming measures in their communities. The work consists of a hardened, low-power traffic sensor, a deployment of a low-cost, wide-area sensor network, a study of optimizing communication efficiency, and a browser-based visualization system for management and analysis. In so doing, we articulated and addressed the following design principles:

[12] http://www.sankey-diagrams.com

- **Devices anywhere**: We designed and deployed a new LP-WAN network supporting sensors anywhere on the road surface. This network exhibits robust coverage, can scale to large deployments, and is supported by measurement and visualization tools for characterizing signal strength.
- **City-friendly integration**: We designed a new family of pavement-mountable sensors with lifetimes comparable to existing, passive RPMs. Installation requires no more expertise than with a traditional RPM. This is accomplished with means for sensor self-localization and sensor-wireles-signal adaptation. Antenna techniques for maximizing signal strength subject to strict packaging considerations were explored in some depth.
- **Open development platform**: We created a flexible software system that allows devices of many sorts to be enrolled in and attached to the network and for geo-referenced data to be displayed graphically. We created flexible data pipelines for visualization and analysis of traffic performance, while further exploring new programming paradigms for creating seamless smart city applications.
- **Edgelessness**: Our programming model supports intelligence in the endpoint sensor itself, and we opened an exploration into the application of on-sensor neural networks for processing magnetometer signals in real time.

Based on the experience in deploying the wireless network and taking field measurements, we have reasoned through the basic challenges and hindrances of deploying traffic sensors in city environments effectively and safely to meet the community-centric goals of quantifying and understanding traffic flows. Three technical papers have been published discussing these issues in some depth.

The study was done in cooperation with the City of Palo Alto's Transportation department that provided important design constraints and success metrics (such as the ability to gather simple sensor data and produce a real-time traffic map). In addition to our ongoing discussions with the City of Palo Alto regarding deployment of our sensors on city streets, we have also opened a discussion with San Mateo County for installing sensors there–both to provide them with real-time data and to support our effort to build a larger dataset of vehicle signatures in support of machine-learning-based vehicle type and speed detection.

In August of this year, we worked closely with IBM Corporation to design, plan and carry out a hackathon that would have included the use of our Traffic Sensors as one of several environmental data feeds. For logistical reasons, IBM chose to cancel the hackathon late in the planning process. As such, the objective of running a hackathon using our sensors remains open. Nevertheless, over the span of this project, two graduate courses at CMU were taught that either contributed to the design of our crowdsourced traffic calming system or used the sensor, network and software in some beneficial way.

V. Recommendations

Our study has yielded a design for city-friendly traffic sensors and tools for management and analytics that even non-specialists could use. It is our goal to eventually deploy these sensors *en masse* and gather traffic data over a wide area.

A significant challenge to this involves finding safe and legal procedures for laying down sensors on active streets. While the devices themselves are not active participants in the flow of traffic, the installation procedures and liabilities for having devices in such potential danger zones limited our capacity to perform actual measurements and experiments.

We recommend the development of policies and frameworks to enable such studies to take place, even perhaps an accessible proving ground for traffic related technologies before continuing to actual city-scale deployments.

Acknowledgements

We wish to thank Josh Mello, former Chief Transportation Official of the City of Palo Alto as well as his team members Rafael Rius and Ruchika Aggarwal for their engagement with us in this project. We also wish to thank Dr. Ray Buettner, Gerald Scott and the staff of the Naval Postgraduate School and the staff of the California Army National Guard at Camp Roberts, California, for making the Joint Interagency Field Experimentation facilities available to us for this study.

References

[1] G. D. Garrod, R. Scarpa, and K. G. Willis, "Estimating the Benefits of Traffic Calming on A Choice Experiment Approach," vol. 36, no. May, pp. 211–231, 2002.

[2] B. Iannucci and A. Rowe, "Crowdsourced Smart Cities," *Intelligent Transportation Systems (ITS) World Congress*, 2017.
[3] LoRa-Alliance, "A technical overview of LoRa and LoRaWAN," no. November, pp. 1–20, 2015. [Online]. Available: https://www.tuv.com/media/corporate/products_1/electronic_components_and_lasers/TUeV_Rheinland_Overview_LoRa_and_LoRaWANtmp.pdf
[4] Semtech Corporation, "LoRa Modulation Basics," Tech. Rep. Revision 2, 2015. [Online]. Available: http://www.semtech.com/images/datasheet/an1200.22.pdf
[5] U. Noreen, A. Bounceur, and L. Clavier, "A study of LoRa low power and wide area network technology," *2017 International Conference on Advanced Technologies for Signal and Image Processing (ATSIP)*, pp. 1–6, 2017. [Online]. Available: http://ieeexplore.ieee.org/document/8075570/
[6] A. J. Wixted, P. Kinnaird, H. Larijani, A. Tait, A. Ahmadinia, and N. Strachan, "Evaluation of LoRa and LoRaWAN for wireless sensor networks," *Proceedings of IEEE Sensors*, vol. 0, pp. 5–7, 2017.
[7] S. Dawson-Haggerty, A. Krioukov, J. Taneja, S. Karandikar, G. Fierro, N. Kitaev, and D. Culler, "BOSS: Building Operating System Services," *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*, pp. 1–15, 2013.
[8] M. Weiss, J. Eidson, C. Barry, D. Broman, L. Goldin, B. Iannucci, E. A. Lee, and K. Stanton, "Time-Aware Applications, Computers, and Communication Systems (TAACCS)," NIST, Tech. Rep., 2015.
[9] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog Computing and Its Role in the Internet of Things," in *MCC '12*, Helsinki, Finland, 2012, pp. 13–15.
[10] M. Satyanarayanan, Z. Chent, K. Hat, W. Hut, W. Richtert, and P. Pillai, "Cloudlets: at the Leading Edge of Mobile-Cloud Convergence (invited paper)," 2014.
[11] D. Astely, E. Dahlman, G. Fodor, S. Parkvall, and J. Sachs, "LTE Release 12 and Beyond," *IEEE Communications Magazine*, vol. 51, no. 7, pp. 154–160, 2013.
[12] J. Gozalvez, "New 3GPP Standard for IoT," *IEEE Vehicular Technology Magazine*, vol. 11, no. 1, pp. 14–20, 2016.
[13] Akamai, "The State of the Internet," Tech. Rep. 3, 2013.
[14] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, "Long-Range Communications in Unlicensed Bands: the Rising Stars in the IoT and Smart City Scenarios," *IEEE Wireless Communications*, pp. 1–7, 2016. [Online]. Available: http://arxiv.org/abs/1510.00620
[15] C. M. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA, USA: Harvard Business School Press, 1997.
[16] S. E. Minson, B. A. Brooks, C. L. Glennie, J. R. Murray, J. O. Langbein, S. Owen, T. Heaton, R. A. Iannucci, and D. Hauser, "Crowdsourced Earthquake Early Warning," *Science Advances*, vol. 1, no. 3, pp. 1–7, 2015. [Online]. Available: http://advances.sciencemag.org/content/1/3/e1500036.full-text.pdf+html
[17] S. Cheung, S. Coleri, B. Dundar, S. Ganesh, C.-W. Tan, and P. Varaiya, "Traffic Measurement and Vehicle Classification with Single Magnetic Sensor," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1917, no. 510, pp. 173–181, 2005. [Online]. Available: http://trrjournalonline.trb.org/doi/10.3141/1917-19
[18] S.-y. Cheung, "Traffic Surveillance by Wireless Sensor Networks : Final Report," *California Path Program Institute of Transportation Studies University of California,Berkeley*, vol. 1, no. January, p. 161, 2007.
[19] A. Haoui, R. Kavaler, and P. Varaiya, "Wireless magnetic sensors for traffic surveillance," *Transportation Research Part C: Emerging Technologies*, vol. 16, no. 3, pp. 294–306, 2008.
[20] A. Tadkase, N. Srinivasan, and B. Iannucci, "Combating Ground Reflection for Wireless Sensors," *Intelligent Transportation Systems (ITS) World Congress*, 2017.
[21] S. Mathur, A. Sankar, P. Prasan, and B. Iannucci, "Energy Analysis of LoRaWAN Technology for Traffic Sensing Applications," *Intelligent Transportation Systems (ITS) World Congress*, 2017.
[22] "Node-RED." [Online]. Available: https://nodered.org
[23] "InfluxData." [Online]. Available: https://www.influxdata.com
[24] "Grafana." [Online]. Available: https://grafana.com
[25] A. Tadkase, N. Srinivasan, and B. Iannucci, "Combating Ground Reflection for Wireless Sensors," in *Intelligent Transportation Systems (ITS) World Congress*, Montreal, Quebec, Canada, 2017.
[26] M. Kong, G. Shin, S. Lee, and I. J. Yoon, "An electrically small, 3D printed folded spherical meander antenna," *2017 Asia-Pacific International Symposium on Electromagnetic Compatibility, APEMC 2017*, pp. 102–104, 2017.
[27] "SUMO - Simulation of Urban Mobility." [Online]. Available: http://sumo.dlr.de