October 29, 2017

Crowdsourced Smart Cities

Robert A Iannucci, *Carnegie Mellon University*
Anthony Rowe, *Carnegie Mellon University*

Paper ID AM-SP1286

Crowdsourced Smart Cities

Bob Iannucci[1], **Anthony Rowe**[2]
1. Carnegie Mellon University, Department of Electrical and Computer Engineering
NASA Ames Research Park, Moffett Field, CA 94035 USA
+1 650 714 1200 – bob@sv.cmu.edu
2. Carnegie Mellon University, Department of Electrical and Computer Engineering, USA

Abstract

The vision of applying computing and communication technologies to enhance life in our cities is fundamentally appealing. Pervasive sensing and computing can alert us to imminent dangers, particularly with respect to the movement of vehicles and pedestrians in and around crowded streets. Signaling systems can integrate knowledge of city-scale traffic congestion. Self-driving vehicles can borrow from and contribute to a city-scale information collaborative. Achieving this vision will require significant coordination among the creators of sensors, actuators, and application-level software systems. Cities will invest in such smart infrastructure if and only if they are convinced that the value can be realized. Investment by technology providers in creation of the infrastructure depends to a large degree on their belief in a broad and ready market. To accelerate innovation, this stalemate must be broken. Borrowing a page from the evolution of the internet, we put forward the notion that an initially minimalist networking infrastructure that is well suited to smart city concepts can break this cycle and empower co-development of both clever city-sensing devices and valuable city-scale applications, with players large and small being empowered in the process. We call this the *crowdsourced smart city* concept. We illustrate the concept via an examination of our ongoing project to crowdsource real-time traffic data, arguing that this can rapidly generalize to many more smart city applications. This exploration motivates study of a number of smart city challenges, crowdsourced or otherwise, leading to a paradigm shift we call *edgeless computing*.

KEYWORDS:

Crowdsourcing, Traffic calming, Smart cities, LP-WAN, LoRa

I. Introduction

With each successive advance in computing and communications technologies, new application frontiers are opened. Complex, real-time signal processing can be done at low power with inexpensive computing devices. Radios have become much more frequency-agile as the vision of the true software-defined radio comes closer to reality. Pervasive sensing and computing could alert us to imminent dangers, particularly with respect to the movement of vehicles and pedestrians in and around crowded streets. Signaling systems could integrate knowledge of city-scale traffic congestion. Self-driving vehicles could borrow from and contribute to a city-scale information collaborative.

These technologies combined with pervasive networking lead us to imagine a world in which everything could be connected. But at present, it isn't. And there are core reasons for this. The vision and puzzle of smart cities compel us to confront the fact that the barriers may not be technological *per se*. Perhaps we need to reconsider the approaches by which we might presume to bring technology to the city.

Cities are living, organic entities with structure, character, core values, momentum and, perhaps most importantly, personalities. Cooper [1], Norman [2] and others have argued that a design-centric approach yields the best match between computing technologies and people and that invisible technologies that just work are preferable to technologies that require people to adapt to them. It stands to reason that the same should be true about cities. When we think about making a city smart, we should probably start with the premise that the city tomorrow will probably work like the city does today with respect to its habits (particularly buying and maintaining infrastructure) and that we should try to make smart city technologies integrate well with these habits.

The smart city concept is nascent, and the interrelationship between cost and value is not well-established. As with other emerging technologies, we might expect a daring few cities to undertake small- to mid-size trials with the outcome being a deeper understanding of the costs and values, rather than an enduring investment in sensors, actuators, radios and computers. The large majority of cities will take a wait-and-see approach. Innovative companies will offer vertical (sensors to applications) solutions in hopes of establishing a first-mover advantage. And smaller technology players–historically better positioned to offer elements or components rather than well-supported vertical services–will wait and watch to see how the model of the smart city evolves, stepping in when niches emerge.

Is this glacial creep fundamental, or might it be possible to accelerate smart city evolution by using a different model of innovation? In this paper, we put forward a new vision that we call *crowdsoruced smart cities.* The vision is founded on the principle that we can and should accelerate the transition from a vertical-solution-only marketplace to one in which at least one key horizontal layer is pre-established: an openly-accessible, wide-area network plus core services, suitable for the types of devices we expect to find in the smart city. This idea borrows from the evolution of the internet itself, recognizing that an initially minimalist networking infrastructure enables innovation by small and large players alike, mitigating to some degree the challenge of obsolescence suffered by many first-movers. This concept aligns with thinking about how the sharing economy could relate to smart cities [3], and seeks to build on experiences from other large-scale smart city experiments such as in Barcelona [4]. Beyond opening a city's infrastructure to citizens, we seek the means to allow anyone to become a data source in the smart city. Rather than viewing the problem as top-down or bottom-up [5], ours can be cast as a middle-out approach.

In Section II, we consider a set of core design principles for the crowdsourced smart city. We highlight some consequences that are a result. In Section III, we reflect on existing cellular networks and the degree to which they do or do not satisfy the design principles. From this, we consider an alternative approach based on low-power wide-area networking (LP-WAN) in Section IV. Armed with a new weapon, the LP-WAN, we consider how the smart city vision might be realized incrementally in Section V. We illustrate the crowdsourcing idea in Section VI by considering a crowdsourced approach to traffic calming. We conclude with reflections and implications.

Contributions: This work presents a vision for accelerating the creation of smart cities and contributes

- An analysis of the dynamics of cities related to adopting would-be smart city solutions.
- A proposal for a network model that helps create incentives for accelerated innovation.
- An example that represents a broad family of crowdsourced smart city applications.
- A perspective on computing platforms that unifies on-device, in-network and in-cloud programming called *edgeless computing*.

II. 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, we do so with some tacit assumptions in mind. The kinds of applications we contemplate depend on (a) the ability to place devices *anywhere* 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, 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 [6]. Rather, we choose to focus on the less-well-understood *out-of-building experience*. This is the domain of smart public spaces, smart transportation systems, self-driving cars, drones, and the like.

A. 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

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

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.

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?

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

We illustrate the importance of this by examining our ongoing project (to be described more fully in Section VI) to establish a network of sensors that can report, in real time, on the traffic conditions in cities. Pittsburgh, Pennsylvania is 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 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 (how frequently is a stop sign maintained?). 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?

C. 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

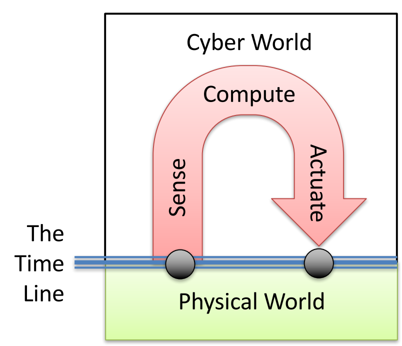

Fig. 1: Simplified representation of information flow in the smart city. Physical quantities (*e.g.,* passing cars, temperature, air quality) in the smart city are sampled by sensor devices. In so doing, their values are associated with the time-of-reading and passed on for processing. Decisions may 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.

[2]http://pittsburghpa.gov/dpw/street-resurfacing

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 [7]. Do the old rules and approaches for third party app developers apply?

D. 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 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 [8]. Satya and colleagues took a complementary view of pushing computation from (at the time) compute-impoverished mobile devices into processors that had cloud-like capabilities but that were proximate to the edge [9]–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, 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?

III. Today's Networks

An immediate and important question is *Do today's 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

[3] http://www.serverwatch.com/news/article.php/1400281

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 [10] and subsequent standards including NB-IoT [11] 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 [12]). 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.

IV. Low-Power Wide-Area Networks

Low-Power Wide-Area Networks (LP-WAN) [13] 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 [14], 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.

A. 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.

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

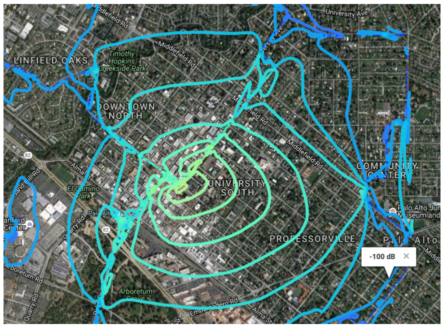

(a) Coverage Map

(b) Gateway

Fig. 2: (a) Map of LoRa coverage around Palo Alto City Hall. Contours of signal strength are shown in increments of 5dB with yellow being the strongest and blue the weakest. Acceptable coverage stretches well beyond the figure's boundaries–in some cases, many miles. In this terrain, a LoRa gateway (b) can cover more than one square mile (image: Multitech).

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 VI.

B. 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.

C. Performance Considerations for LP-WAN

While bandwidth for LP-WANs is in no way comparable to cellular, the coverage possible with a single gateway, and hence the cost per unit area to deploy a network, is remarkable. In our own research, we have built test networks covering portions of the City of Palo Alto, California and the Oakland neighborhood of Pittsburgh, Pennsylvania.

Figure 2 depicts the coverage from a single test LoRa gateway mounted on the roof of Palo Alto's City Hall. Signal strength data was collected with a combination of drive- and walk-testing. Most areas within a one mile radius of City Hall have coverage adequate for many smart city applications. Notably, this signal was able to reach many miles beyond. The cost was a single gateway connected by a Power-over-Ethernet wire to an existing wired network.

D. 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...

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.

V. Crowdsourcing 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.

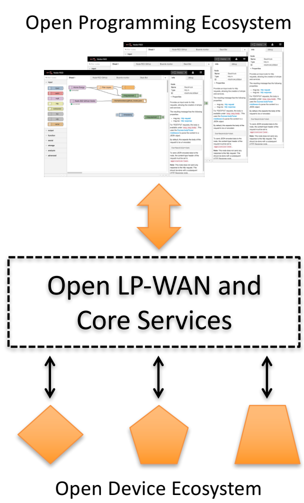

Fig. 3: An open LP-WAN with core services and edgeless computing is the essential "middle" of the crowdsourced smart city.

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 makers, high school science students and application developers with the challenge to build novel smart city apps. While at first blush this may seem outlandish, the possibilities are quite real.

In our work, we are creating reference implementations of LP-WAN-enabled small circuit boards that can be easily integrated with sensors to perform a wide range of sensing functions (we examine one such in the next section). 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 are following this trend by creating the basic boards and applications out of which many projects might grow. Concurrently, we are working with forward-looking city partners in Palo Alto and Pittsburgh to deploy initial LP-WANs to be operated as proofs-of-concept. Our objective is to open this smart city infrastructure to hardware and software developers

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

and to encourage wide participation in an incremental build-out of two smart cities. It is our belief that the basic step of building an LP-WAN and a few supporting web services may be sufficient to incite a kind of innovation revolution. We call this 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.

The middle-out concept is depicted in Figure 3. By seeding a city with a basic LP-WAN and complementing the raw network with core services (Figure 4), we enable both the simple attachment of a wide range of new sensor types to the city and a corresponding open programming ecosystem. 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 to make data available to one another, we potentially open new kinds of applications that no one might have been able to undertake.

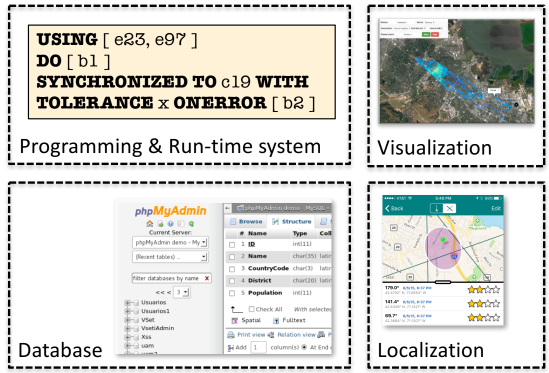

Fig. 4: Minimal LP-WAN services to support crowd-sourced smart cities.

This level of enablement not only makes city-sponsored projects possible, it makes accessing and processing city data accessible to a wide range of non-specialists. We imagine city-level information hackathons and high school science projects that create new sensor types and mash up data from many city sources.

We further imagine that the power of the crowd can be leveraged to tie smartphones into this city-wide network. With suitably-programmed Bluetooth, WiFi or other radios in these circuit boards (such as we are building) and apps in the phones, it is possible to consider using phones as "data mules" to bulk-transfer information from sensors to phones and, thereby, to our core network services. In reverse, phones could provide a channel by which larger software updates are delivered to theses sensors. Bluetooth beacons in the sensors could be used by self-driving (and other) cars as local information relay points. Similarly, beacons in phones could be detected by the sensors to provide personalized information services.

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. These are important research vectors that the crowdsourced smart city vision can fully open.

VI. Example

To make the concepts of crowdsourced smart cities more concrete (so to speak), and to emphasize the value of investing in LP-WAN networks as enablers in the middle of the smart city's infrastructure, we have undertaken an application that presumes such a middle layer and builds on it with some simple sensors *at the bottom* that wirelessly connect, and some analytic tools *at the top* that access the data.

Our example focuses on creating a network of inexpensive, easily-installed sensors to map traffic flows accurately. The reader is encouraged to reflect on this as an example of a potentially much broader class of of smart city applications that could build on the base of an LP-WAN *middle*.

A. The Challenge

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 [17].

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 amount to just a small number of pneumatic-hose vehicle sensors per city. As a result, the information collected is limited, hampering rigorous quantitative analysis of traffic flows with and without specific calming interventions.

B. A Novel Solution

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.

In partnership with the City of Palo Alto, we have launched a project, *Crowdsourced Traffic Calming*, to address this challenge. 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.

Our sensor board is small (3" x 2.5") and 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.

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.

C. The Sensor

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

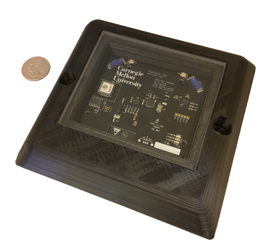

Fig. 5: CMU's TrafficDot – an intelligent RPM for counting vehicles and measuring speed.

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

In order to correctly geo-reference its readings, the location of each TrafficDot must be recorded. This can be accomplished (a) at installation time by manually recording each TrafficDot'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 TrafficDot itself (we've built in a GPS receiver for this purpose). We are exploring the accuracy by which the network can localize TrafficDots. Network-based localization, if adequately accurate, could make geo-referencing of TrafficDots transparent to city practices and could obviate inclusion of a GPS module.

[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

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

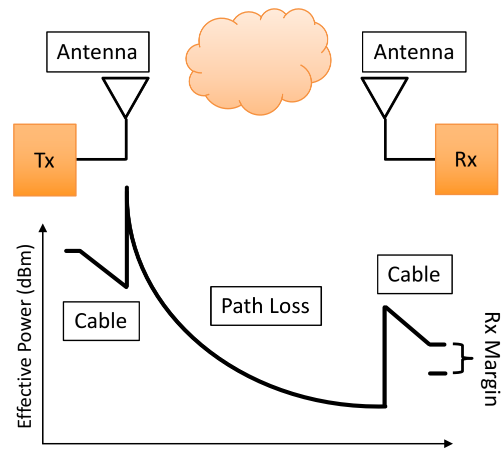

Fig. 6: Signal strength at the receiver is a function of transmitter signal strength, cable losses, antenna gains (or losses), and free space path loss. With power being capped by regulation and receive sensitivity set by technology, maximizing distance is done by optimizing the antenna subsystems.

D. 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:

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

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 [18].

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.

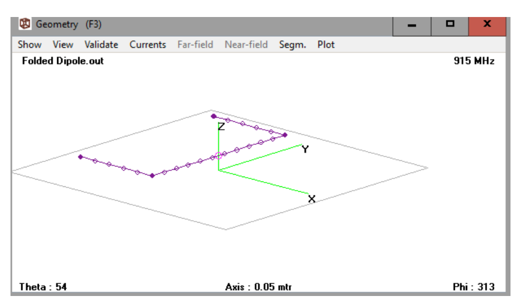

Fig. 7: A 915 MHz dipole antenna, folded to fit inside an RPM.

Ground proximity: A second problem is the effect the pavement itself has on the antenna's pattern. Mounted in an RPM, the 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 takeoff 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.

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 TrafficDot 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 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° takeoff 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.

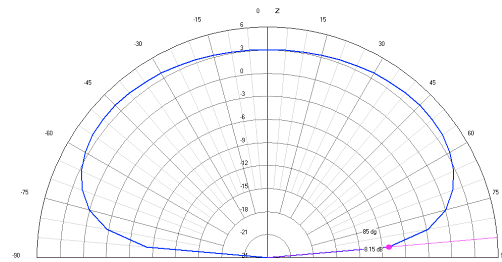

Fig. 8: Geometries and placement of smart city devices can compromise antenna design. In this case, a folded dipole antenna packed into a small sensor mounted on pavement would exhibit near-vertical-incidence behavior (main lobe points upward). Gain at a low takeoff angle (in red)–such as would be the 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. [18]

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 takeoff angle demanded of our TrafficDots and provides an adaptive means for the TrafficDot to beam-steer its signal, using a combination of gateway signal strength measurement and a beam-forming antenna array. A deeper treatment of our approach to antenna optimization is given in our companion paper [18].

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 TrafficDot battery lifetime in another companion paper [19].

E. Network and Analytics

In our experiments, we are building out the Palo Alto LoRa network incrementally. We have the capability to establish new gateways quickly, and we’ve also created a mobile RF laboratory both for studying placement of network elements and for creating hastily-formed networks. The mobile laboratory has a software-defined infrastructure together with stand-alone networking, computing and power infrastructure as well as a variety of internet backhaul capabilities, including a satellite link that is more than adequate for LP-WAN applications.

While a purpose-build sensor network will have a full stack of cloud-side software created for a specific purpose, we are creating a minimal set of basic network services common to smart city applications. A common geo-indexed database serves as a data repository. A programming environment will support creation of city-scale apps. Localization services bind TrafficDots to physical locations. Basic visualization tools support extraction of geo-referenced data from the database and creation of map overlays. Specific applications, such as traffic calming, can build on this base to generate heatmaps, Sankey diagrams,[10] or other presentations as appropriate. It is our hope that an open visualization toolkit will encourage creation of additional openly-available layers.

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

VII. Conclusion

We have outlined core design principles to consider when selecting a networking infrastructure for smart cities. Notions of devices anywhere, city-friendly integration, an open development platform, and edgelessness are, we argue, necessary conditions for igniting an innovation revolution for the smart city. We have explored the relative merits of existing cellular networks and LP-WANs for this purpose and have shown that LP-WANs may be an important disruptive technology to explore and exploit in the pursuit of the smart city.

We have argued that LP-WANs together with the power of the crowd can lead to crowdsourced smart cities. We have opened technology explorations for a few of the key challenges. In addition, we have outlined a few re-usable services that complement a middle-out smart city design.

We are applying these techniques and concepts together with the City of Palo Alto toward real-time, crowdsourced traffic calming. It is our hope that these experiments and the community that might adopt some or all of our techniques will help bring about a crowdsourced smart city revolution.

VIII. Acknowledgements

This work was supported in part by the National USDOT University Transportation Center for Safety for which the authors are grateful. We wish to thank Joshua Mello, AICP, Chief Transportation Official of the City of Palo Alto for bringing the traffic calming challenge to our attention and for engaging with us in the creation of this novel crowdsourced traffic calming prototype. We also wish to thank Director Ken Dueker and Simon Williams of the Office of Emergency Services, City of Palo Alto, for supporting our efforts to establish an experimental LP-WAN network on which these studies have been based. We are also grateful for gifts in support of our research from Cisco, CableLabs, the City of Palo Alto, Verizon, CrownCastle, and Symmetry Electronics. Thanks go to students in the Wireless Systems concentration at CMU Silicon Valley and especially to Ervin Teng, João Diogo Falcão, Cef Ramirez and Ashutosh Tadkase for their contributions to our visualization, database, signal strength analytics and microservices infrastructure as well as the TrafficDot and its intellectual predecessor, the PowerDué.

References

[1] A. Cooper, *The Inmates Are Running the Asylum*. Indianapolis, IN, USA: Macmillan Publishing Co., Inc., 1999.

[2] D. A. Norman, *The Invisible Computer*. Cambridge, MA, USA: MIT Press, 1998.

[3] E. Almirall, J. Wareham, C. Ratti, P. Conesa, F. Bria, A. Gaviria, and A. Edmondson, "Smart Cities at the Crossroads," *California Management Review*, vol. 59, no. 1, pp. 141–152, 2016. [Online]. Available: http://journals.sagepub.com/doi/10.1177/0008125616683949

[4] P. Berrone, J. E. Ricart, and C. Carrasco, "The Open Kimono," *California Management Review*, vol. 59, no. 1, pp. 39–70, 2016. [Online]. Available: http://journals.sagepub.com/doi/10.1177/0008125616683703

[5] B. Cohen, E. Almirall, and H. Chesbrough, "The City as a Lab," *California Management Review*, vol. 59, no. 25, pp. 5–13, 2016.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] J. Gozalvez, "New 3GPP Standard for IoT," *IEEE Vehicular Technology Magazine*, vol. 11, no. 1, pp. 14–20, 2016.

[12] Akamai, "The State of the Internet," Tech. Rep. 3, 2013.

[13] 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, 2015. [Online]. Available: http://arxiv.org/abs/1510.00620

[14] Semtech, "LoRa Modulation Basics," no. May, pp. 1–26, 2015. [Online]. Available: http://www.semtech.com/images/datasheet/an1200.22.pdf

[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] 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.

[18] A. Tadkase, N. Srinivasan, and B. Iannucci, "Combating Ground Reflection for Wireless Sensors," in *ITS World Congress*, Montreal, Quebec, Canada, 2017.

[19] S. Mathur, A. Sankar, P. Prasan, and B. Iannucci, "Energy Analysis of LoRaWAN Technology for Traffic Sensing Applications," in *ITS World Congress*, Montreal, Quebec, Canada, 2017.