NORTHWESTERN
UNIVERSITY

Electrical Engineering and Computer Science Department

Technical Report
NWU-EECS-09-10
April 8, 2009

C3R – Participatory Urban Monitoring from your Car

John S. Otto, John P. Rula, and Fabián E. Bustamante

Abstract

Detailed on-line measurements of city air quality are necessary to support compliance with air quality standards and emission strategy developments and to inform both policy makers and the general public in a timely manner. Current approaches to pollution monitoring, however, rely on very small numbers of stationary nodes. Providing detailed air quality information is a particularly hard challenge given the explosive and mostly uncontrolled growth of urban populations and the high cost of typically employed monitoring solutions.

We present *C3R*, an alternative, participatory approach to urban air quality monitoring. *C3R* takes advantage of recent advances in vehicular ad-hoc networks and sensing technology. Vehicular networks provide an ideal platform for urban monitoring, with hundreds or thousands of potential nodes, freed from space and energy constraints, and reaching nearly every corner of modern metropolitan areas. *C3R* nodes, each equipped with low-cost pollution sensors and more standard location and communication devices, collect and disseminate pollution level information to other vehicles through a basic gossiping protocol. Each node maintains a real-time detailed air quality map that is shared with the driver, other vehicles and public agencies. We describe our current implementation of *C3R* and discuss some of the challenges we have addressed, including the adaption of solid-state pollution sensors to provide accurate measurements under dynamic vehicular conditions.

Keywords: Application, Environmental pollution, Vehicular networks, Systems

C3R – Participatory Urban Pollution Monitoring from your Car

John S. Otto John P. Rula Fabián E. Bustamante
Department of Electrical Engineering and Computer Science
Northwestern University
{jotto,john.rula,fabianb}@eecs.northwestern.edu

Abstract

Detailed on-line measurements of city air quality are necessary to support compliance with air quality standards and emission strategy developments and to inform both policy makers and the general public in a timely manner. Current approaches to pollution monitoring, however, rely on very small numbers of stationary nodes. Providing detailed air quality information is a particularly hard challenge given the explosive and mostly uncontrolled growth of urban populations and the high cost of typically employed monitoring solutions.

We present *C3R*, an alternative, participatory approach to urban air quality monitoring. *C3R* takes advantage of recent advances in vehicular ad-hoc networks and sensing technology. Vehicular networks provide an ideal platform for urban monitoring, with hundreds or thousands of potential nodes, freed from space and energy constraints, and reaching nearly every corner of modern metropolitan areas. *C3R* nodes, each equipped with low-cost pollution sensors and more standard location and communication devices, collect and disseminate pollution level information to other vehicles through a basic gossiping protocol. Each node maintains a real-time detailed air quality map that is shared with the driver, other vehicles and public agencies. We describe our current implementation of *C3R* and discuss some of the challenges we have addressed, including the adaption of solid-state pollution sensors to provide accurate measurements under dynamic vehicular conditions.[1]

1 Introduction

Air pollution, including byproducts of industrial processes and automotive emissions, is directly responsible for a number of respiratory infections, asthma, heart disease and lung cancer. According to the World Health Organization, environmental risk factors contribute to 24% of the global burden of disease from all causes and to 23% of all deaths [9, 10].

[1]Currently under submission.

Detailed on-line measurements of city air quality are necessary not only to support compliance with standards and emission strategy developments, but to inform both policy makers and the general public in a timely manner. Current approaches to pollution monitoring, however, rely on very small numbers of stationary nodes. For instance, the 2009 Illinois air monitoring network includes fewer than 20 monitoring sites for Chicago's Cook County – the second most populous county in the United States and the worst in terms of dangerous air pollution [8].

Providing detailed air quality information is a particularly hard challenge given the explosive and mostly uncontrolled growth of urban populations and the high cost of typically employed monitors. A carbon monoxide sensor costs nearly $20000, and particulate matter sensors can cost as much as $30000 [2, 6]. As a city's population grows, the number of sensor sites needed to ensure appropriate coverage quickly becomes prohibitively expensive.

In this paper we present *C3R*, an alternative, participatory approach [12] to urban air quality monitoring. *C3R* takes advantage of recent advances in vehicular ad-hoc networks and sensing technology. Vehicular networks provide an ideal platform for urban monitoring, with hundreds or thousands of potential nodes, freed from space and energy constraints, and reaching nearly every corner of modern metropolitan areas. *C3R* nodes, each equipped with low-cost pollution sensors and more standard location and communication devices, collect and disseminate pollution level information to other vehicles through a basic gossiping protocol. Each node maintains a real-time detailed air-quality map that is shared with the driver, other vehicles, and public agencies.

Our latest realization of a *C3R* node incorporates a low-cost ($<$\$50) sensor apparatus for monitoring ambient temperature, humidity and carbon monoxide concentrations. Carbon monoxide (CO), a product of incomplete burning of hydrocarbon-based fuels, is a colorless and odorless but seriously toxic gas, and is particularly poisonous for infants, elderly persons and individuals with respiratory diseases. We limit the cost of our sensor apparatus by adapting a Hanwei Electronics MQ-7 CO sensor [7] (with a per unit cost of $9) for use at high-speed – a condition

for which it was not originally designed.

Using a set of these nodes, we survey the distribution of CO concentration in downtown Chicago and show significant spatial variation with high concentration gradients (5 ppm over distances of less than a 100 m block), that traditional, sparse and stationary monitoring stations would fail to capture. We use our experimental observations to evaluate the possibility of a participatory, VANET-based approach for building detailed maps of pollution concentration. We do this through simulation, employing a basic, well-understood gossiping communication protocol [31], together with vehicular mobility [14] and signal propagation [28] models. Our results show that using a simple gossiping protocol, a new report can spread quickly through the network, reaching 90% of the nodes in an 8 km^2 area in 12 minutes. This allows each node to quickly assemble an up-to-date pollution map that contains the majority of recently collected messages – nearly 50% of the measurements collected in the previous 15 minutes. The map gives a picture of pollution that is predominately composed of recent reports from within 1.0 km of the node's location–the average distance from the node to the pollution measurement's position – but also contains measurements from up to 2.4 km away. As a result, the map provides local, relevant pollution measurements to the driver, in contrast with a stationary sensor that may not sample the air near the node.

The paper makes the following main contributions:

- A description of *C3R*, a new participatory approach to urban pollution monitoring, and its realization. *C3R* is the first (to our knowledge) mobile end-to-end system for monitoring and reporting environmental pollution concentrations.
- The design, implementation and evaluation of a low-cost, mobile pollution monitoring platform.
- An analysis of the spatial distribution of urban pollution showing the need for a distributed, mobile approach to monitoring.
- Results from a detailed simulation study showing the timeliness and coverage of pollution measurements that can be distributed through a basic VANET gossiping protocol.

The rest of the paper is structured as follows. After a review of background and closely related work in Section 2, we describe the *C3R* approach for urban environmental monitoring in Section 3. Section 4 discusses the necessary adaptations to allow accurate pollution sensing from a moving vehicle. We examine the spatial distribution of urban air pollution and demonstrate the feasibility of *C3R* to provide a real-time map of pollution to drivers in Section 5. Finally, we discuss future work and conclude in Section 6.

2 Background

The significant health risks associated with elevated pollution levels are strong motivations for air pollution monitoring systems. Measurements can signal pollutants reaching hazardous levels in a city, allowing government to enact targeted programs to reduce the concentrations of those most problematic pollutants. Monitoring systems are able to reveal long term trends in pollution, and current measurements can be compared to previous conditions or quantified with metrics such as the Air Quality Index [1].

Current monitoring systems with stationary sensors are limited in their ability to observe the actual spatial distribution of air pollution, and as a result are unable to identify regions where air quality may be particularly poor [8]. Fine-grained measurement studies of urban pollution have revealed significant variation in air quality (e.g. carbon monoxide (CO) concentration) from street to street [16], and indicated some of the factors that determine air pollution levels such as freight transports, urban canyons and wind direction [32]. Lena *et al*. [24], for instance, correlated increased heavy truck traffic with higher levels of air pollution in the Hunts Peak neighborhood of New York City.

It is important to identify localized peaks of pollution because people in these areas may be subject to a higher risk of disease. Without adequate monitoring spatial resolution, people may not be aware of their increased risk factors. The results of the Hunts Peak study estimated that some residents were exposed to pollution levels exceeding thresholds that will cause respiratory problems [24]. Simple changes in freight routing can impact children health, as it has been shown that children attending schools near high truck traffic experienced more respiratory symptoms relative to children at schools near low truck traffic [22]. A measurement system with a small number of stationary nodes lacks the spatial resolution to detect these localized peaks of pollution. If greater resolution were attained, public officials could be alerted of the problem and enact programs to address the pollution's source.

Our work aims at enabling distributed, live pollutant monitoring by leveraging advancements in sensing technology and vehicular network data dissemination protocols. Low-cost, high-quality devices for sensing, storage, and communication make wide adoption possible. The high mobility of vehicular nodes results in an ever-changing set of vantage points for measuring urban pollution, as well as a highly dynamic connectivity graph that ensures rapid dissemination of measurements.

Several projects have studied environmental monitoring from mobile devices (e.g. [5, 17, 18, 21, 23, 25, 26, 29, 33]). MIT's CarTel project uses a deployment of instrumented vehicles to measure traffic conditions in the road network, collecting data in a centralized database

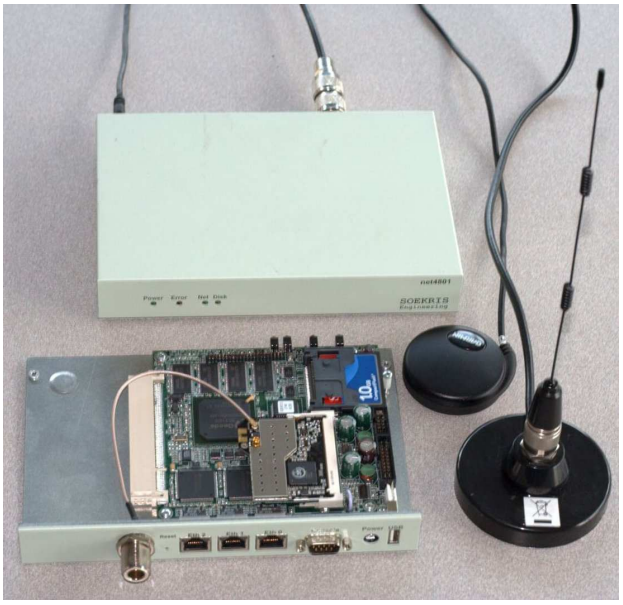

(a) *C3R* Soekris node (left) with Garmin GPS 18 (middle) and 7 dBi 802.11 antenna (right)

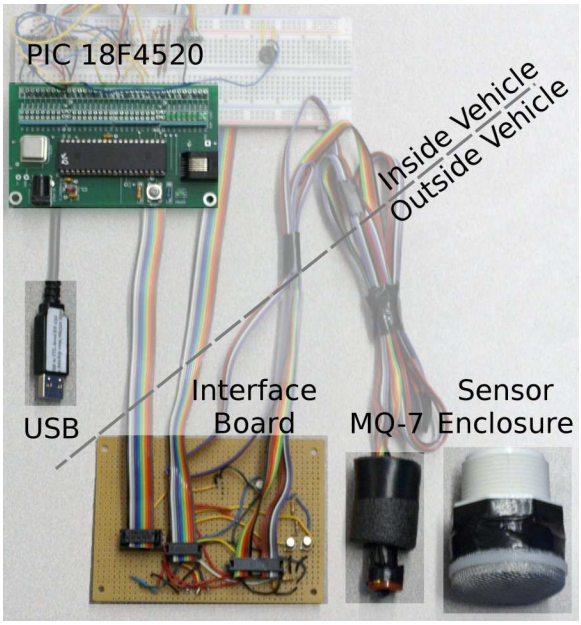

(b) *C3R* sensing components with PIC 18F4520 microcontroller (top left), USB connection (left), Interface Board (bottom left), Hanwei MQ-7 CO sensor (bottom middle) and Insulating Sensor Enclosure (bottom right)

Figure 1: *C3R* Node and Sensing Hardware

via opportunistic WiFi connections [21]. Instead of instrumenting vehicles, Nericell relies on mobile smartphones to infer road and traffic conditions [25]. Targeting the monitoring of infrastructure condition, the Pothole Patrol uses instrumented vehicles to identify areas with poor road surface, and tags their locations in a centralized database [18]. Close in spirit, Walkability [29] aggregates geotagged photos of sidewalk cracks submitted by users.

In the context of pollution monitoring, Ghanem *et al.* [19] proposes the use of sensor grids of stationary nodes for distributed sensing and discusses approaches to the analysis and visualization of pollution data, while Croxford *et al.* [16] examines the spatial diversity of urban pollution measurements in relation to the city's street grid using a deployment of stationary roadside pollution sensor nodes. BikeNet collects data from bicycles instrumented with sensors that monitor environmental parameters (carbon dioxide–CO_2) in addition to the rider's physiological statistics while riding [17]. North *et al.* [27] presents the MESSAGE architecture for sensing multiple pollutants from vehicular nodes over ad-hoc networks using a novel spectroscopic UV sensor. For semiconductor sensors, drift over time of the sensitivity to CO is a known issue. An alternative to periodic recalibration is proposed by [30], in which vehicles could recalibrate their sensors against each others' measurements. The authors show that periodically recalibrating against another vehicle's measurement eliminates drift trends. In contrast, *C3R* adopts a participatory approach for mobile sensing building on adapted, low-cost pollution sensors.

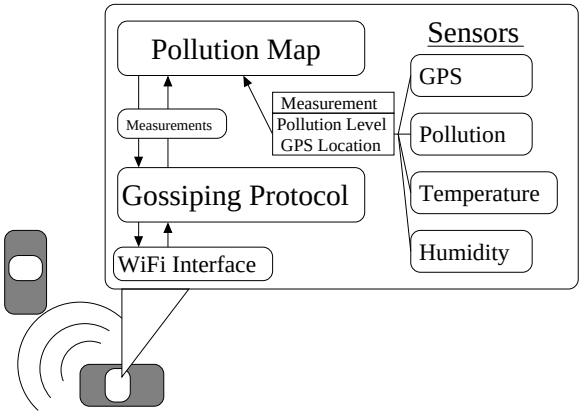

Figure 2: *C3R* Architecture. GPS and sensors are used to periodically take pollution readings, which are added to the map. The gossiping protocol exchanges measurements with other nodes to propagate measurements and build up its own map of pollution.

3 *C3R* Architecture

This section presents a detailed description of *C3R*, our vehicular-based, participatory approach to urban pollution monitoring, and its current realization.

The detailed on-line monitoring of air quality is a problem that fundamentally requires a mobile sensing solution. The observed variations in pollution concentration levels across a region, the mostly uncontrolled growth of urban environments, and the associated cost and logistic complexities make infeasible any solution based on static monitoring stations. *C3R* is motivated by the observation that vehicular networks can provide an ideal platform for urban mobile sensing [21], virtually free of space and power constraints and covering nearly every corner of the city.

Each participating *C3R* vehicle is equipped with low-cost pollution sensors, in addition to more common location, communication and computational devices. Each node monitors the environment, basically as a side-effect of normal vehicular operation, and exchanges the collected information with other participating vehicles through a basic epidemic protocol [31]. Epidemic approaches for data distribution make minimal assumptions about the connectivity of the underlying network, relying instead on communication between random pairs of mobile nodes to distribute messages across the connected parts of the network. The aggregated information can be used to generate a detailed map of air quality that is made directly available to the driver, and/or provided to third parties (e.g. environmental monitoring agencies) without compromising drivers' privacy [20].

To ensure wide adoption, we designed *C3R* monitoring platform (Figure 2) to be easy to build and very economical. Nodes require a GPS device for geolocation and a WiFi radio to communicate with other nodes. A microprocessor takes readings from the pollution, temperature, and humidity sensors and reports the data to the node. Our current implementation costs $580 for the node and sensors, and includes a fanless Soekris net4801 system ($300), a Garmin GPS 18 device ($100), and an Ubiquiti Networks SuperRange2 WiFi radio ($100) with a Pacific Wireless 7dBi 802.11 antenna ($30) for communication (Figure 1a). The sensing apparatus (Figure 1b), costing approximately $50 in total, is built around a Microchip PIC 18F4520 microcontroller ($8). We connect to the PIC a Hanwei Electronics MQ-7 CO sensor [7] ($9), a temperature sensor ($2), and a humidity sensor ($15). We implement a well-understood gossiping protocol [31] over which the nodes exchange their measurements. As nodes communicate with each other, they build up their own map of recent pollution measurements.

4 *C3R* Pollution Data Acquisition

Although our approach is general, we focus this study on the monitoring of a particular pollutant – carbon monoxide (CO). CO is a product of incomplete burning of hydrocarbon-based fuels. This colorless and odorless gas is highly toxic, especially to infants, elderly persons and individuals with respiratory diseases. Each participating *C3R* vehicle is equipped with a low-cost CO sensor that is used to passively monitor the environment as aa side-effect of the vehicle's normal operation. Since vehicles are a significant source of carbon monoxide – 40% of all CO emissions in Cook County [11] – measuring from the primary source of the pollution will capture the strong spatial correlation between vehicles and elevated CO levels.

There is a variety of high-end air pollution sensors that provide highly accurate and precise measurements [15]. Their costs, however, make them unsuitable for large-scale deployment. Non-Dispersive InfraRed (NDIR) sensors, for instance, are general and highly accurate devices with per unit cost of approximately $600 dollars. To ensure wide adoption for *C3R*, we opt instead for low-cost pollution sensors such as electrochemical and solid-state devices. Electrochemical sensors are relatively inexpensive and offer sufficient sensitivity in a small form factor. However, they rely on a depletable reactive element, and they may have to be replaced more than than once a year. We selected solid-state semiconductor sensors as they offer a lifespan of years, are sufficiently sensitive to detect typical urban CO levels, and are extremely affordable.

One of the main challenges of working with low-cost sensors, in general, is that while most of them are effective in the controlled and fairly static settings for which they were designed, their responses to uncontrolled scenarios are unknown. The following paragraphs introduce the sensors we use for *C3R*, the process of calibration, and the modifications we explored and implemented to ensure their correct operation in the context of vehicular networks.

4.1 Sensor Operation

We use the Hanwei Electronics MQ-7 CO sensor [7], shown in Figure 3, which has a tin dioxide (SnO_2) semiconductor element that reacts with CO. As the concentration of CO increases, the resistance of the SnO_2 material decreases. The right side of Figure 3 shows the semiconductor element (the cylinder) and a built-in heating coil (passing through the cylinder).

The heater alternates between high and low temperature, switching the sensor between its "purge" and measurement cycles. The "purge" cycle at high heat frees the CO that accumulated on the sensor in the previous measurement, which serves to reset the sensor for the next measurement cycle. The sensor's specifications require that a 60 second "purge" cycle at high heat (5 V across a 33 Ω resistor dissipates 760 mW of power) precedes every 90 second measurement cycle at low heat (1.4 V, 33 Ω resistor, 60 mW of power).

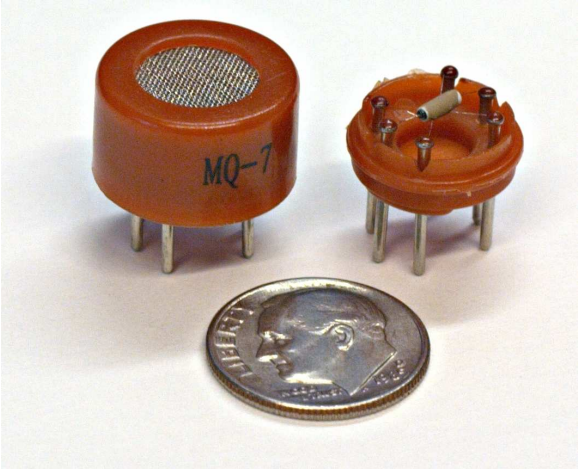

Figure 3: MQ-7 Sensor, with external housing removed on right.

To drive this purge-measure heating cycle, perform analog to digital conversions, and report measurements to the node via a serial data connection, we program a Microchip PIC 18F4520 microcontroller. In addition to the CO sensor, we also utilize temperature and humidity sensors to monitor ambient and sensor conditions. Figure 1 shows all the components of our sensing apparatus, as well as the Soekris embedded system that manages data collection and provides location and communication capabilities via a GPS device and IEEE 802.11b radio.

4.2 Sensor Calibration

To interpret data from low-cost sensors, it is necessary to apply a series of conversions to translate from the low-level measurement (e.g. sensor resistance) to a useful metric. Gas concentrations are typically measured in parts per million (ppm), which is defined as "a volume over volume ratio which expresses the volumetric concentration of a gaseous air contaminant in million unit volumes of gas" [3]–so, 1 ppm CO means that there is one unit volume of CO for every million unit volumes of air. To obtain concentration measurements, we first examine the sensor circuit to determine how to measure sensor resistance. Then, we account for variation in sensor output due to temperature and humidity factors. Finally, we show how we calibrate the CO sensor, which ultimately yields a model that translates sensor resistance to CO concentra-

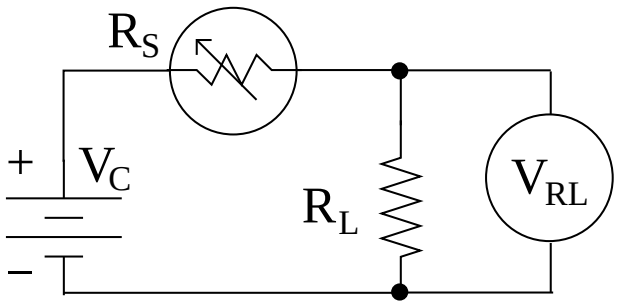

Figure 4: Sensor measurement circuit. R_S is the sensor; its resistance varies with CO concentration. By measuring the voltage drop over a resistor (V_{R_L}) with known resistance ($R_L = 10\text{k}\Omega$), we can determine the resistance of the sensor.

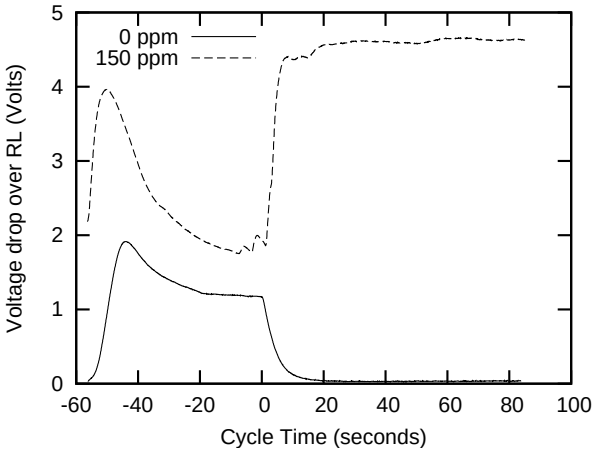

Figure 5: Measured V_{R_L}. Higher CO concentrations result in higher V_{R_L} measurements during the measurement phase (Cycle Time > 0).

tion.

The sensor parameter we wish to measure is resistance, but this is difficult to measure directly. We infer sensor resistance (R_S) by measuring the voltage drop over a resistor (V_{R_L}) with known resistance ($R_L = 10$ kΩ) that is in series with the sensor (circuit shown in Figure 4). Analog-to-digital conversion takes a voltage (in this case, V_{R_L}) and returns an n-bit integer that represents a fraction of the circuit voltage ($V_C = 5$ V). With this data, we can calculate R_S given V_{R_L} [7]:

$$\frac{R_S}{R_L} = \frac{V_C - V_{R_L}}{V_{R_L}}$$

Figures 5 and 6 show normal V_{R_L} and R_S values over the course of a purge-measurement cycle in clean air and 150 ppm CO. The x-axis is the time in the cycle, where zero is the transition from the purge phase to the measure-

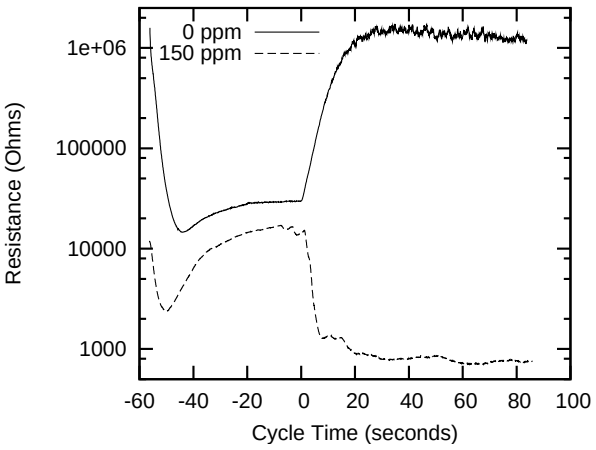

Figure 6: Sensor Resistance. Higher CO concentrations result in lower sensor resistance values during the measurement phase (Cycle Time > 0).

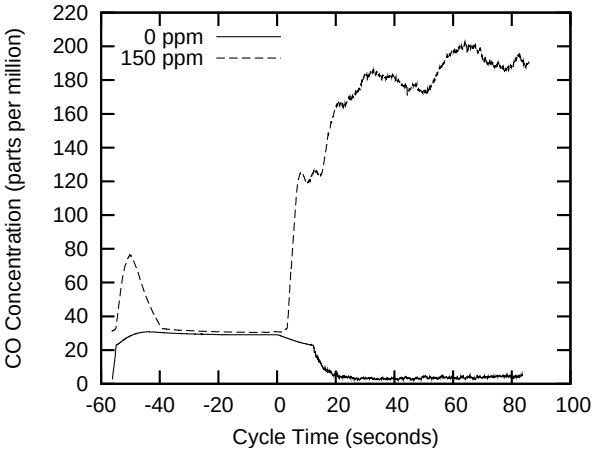

Figure 7: Sensor CO Concentration. Using our empirically-determined calibration data, we convert resistance values to CO concentration (ppm).

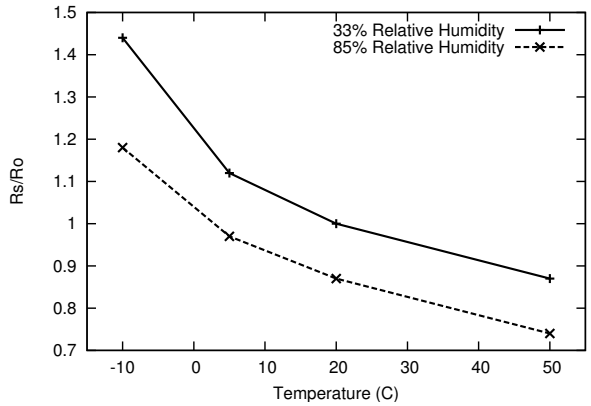

Figure 8: MQ-7 sensor resistance correction factors [7] for temperature and humidity. $\frac{R_S}{R_O}$ is a ratio between measured sensor resistance and R_O, which is the reference sensor resistance at 100 ppm CO. As temperature and humidity increase, the sensor underestimates the resistance and overestimates the resulting CO concentration.

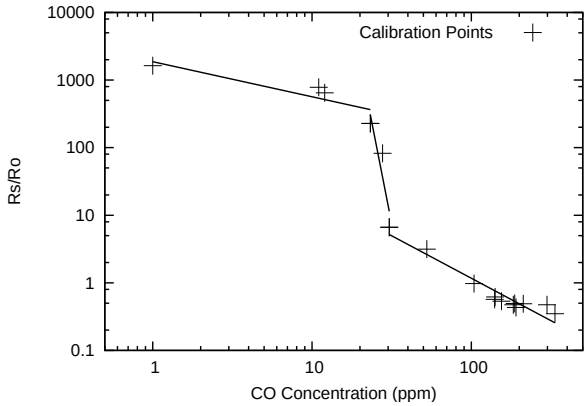

Figure 9: Sensor calibration curve, for converting from R_S to CO concentration (ppm). The data points are recorded sensor resistances under a given CO concentration reported by our ground truth sensor. We use three linear regressions to approximate the relationship between sensor resistance and CO concentration.

ment phase. There is an inverse relationship between V_{R_L} and R_S; as the voltage drop over R_L decreases, R_S increases. The sensor's resistance decreases with greater air temperature and relative humidity (as shown in Figure 8); we incorporate this correction factor into our resistance measurements.

Accounting for the temperature and humidity corrections standardizes the outputs from the sensor, but a more fundamental property of the sensor varies from sensor to sensor: reference resistance at 100 ppm CO. We observed R_O values from 1420 to 2760 Ω for the sensors that we calibrated. This means that a simple initial calibration procedure must be performed on each sensor. The R_O value (e.g. in the $\frac{R_S}{R_O}$ ratio on the y-axis in Figure 8) is used to standardize the measured sensor resistances for use in computing CO concentration.

The MQ-7 datasheet specified the relationship between sensor resistance and CO concentration for 50 to 4000 ppm, but we are most interested in measuring typical urban CO concentrations that are lower than 50 ppm. Since we can not assume that the same relationship holds at concentrations < 50 ppm, we collect sensor resistances over a broad range of CO concentrations. To obtain accurate measurements of CO concentration, we use an Enmet RECON CO meter [4], a portable CO monitor that provides measurements of CO at 1 ppm precision with as response time of a few seconds to changes in CO concentration. Since these devices are calibrated at the factory and are primarily intended for personal safety applications, we consider them to be a reasonable ground truth source of CO concentration.

In our calibration experiment, we measure CO concentrations from 11 to over 300 ppm. We collect our calibration measurements in a subterranean parking structure (that had minimal air flow). Using the exhaust of a vehicle as a CO source, we place our sensor apparatus next to the RECON meter, and vary the distance between the exhaust pipe and the sensors. We added an additional data point from clean, outdoor air, and plot this value at 1 ppm. The MQ-7 datasheet specifies that there should be a linear relationship on a log-log scale of CO concentration versus $\frac{R_S}{R_O}$; Figure 9 shows the calibration curve that we recorded. While our data is approximately linear for the range specified in the datasheet, the relationship changes significantly at concentrations less than 30 ppm CO. We observe a similar calibration curve for all sensors that we used, suggesting that this curve could be generalized to allow for simple calibration of new sensors with a small number of calibration measurements. To model the relationship between sensor resistance and CO concentration, we compute linear regressions over three intervals of CO concentration (x-axis) and map sensor resistance measurements to CO concentration using these approximation lines. Figure 7 shows the CO concentration values computed by our fitted model in clean air and 150 ppm CO.

4.3 Sensing from a Vehicle

We observe that changing the environmental conditions–including variations in temperature and wind speed–has unexpected impacts on the sensor's output. Using knowledge of how the sensor operates, we first show how we identify that the inaccuracy stems from insufficient oper-

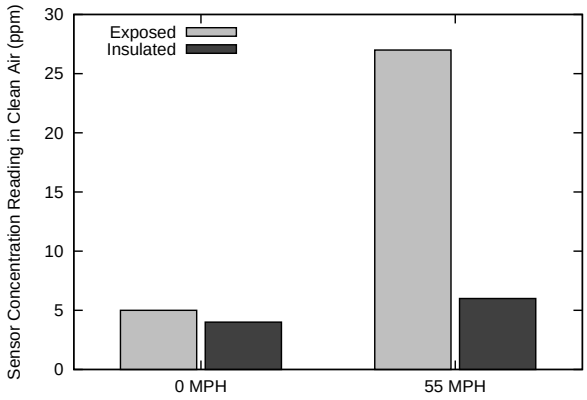

Figure 10: Error between actual CO concentration and the measured CO concentration derived from the sensor's resistance measurements. The exposed sensor has significantly higher error when traveling at high speed, in comparison to the insulated sensor that is able to function consistently under both conditions.

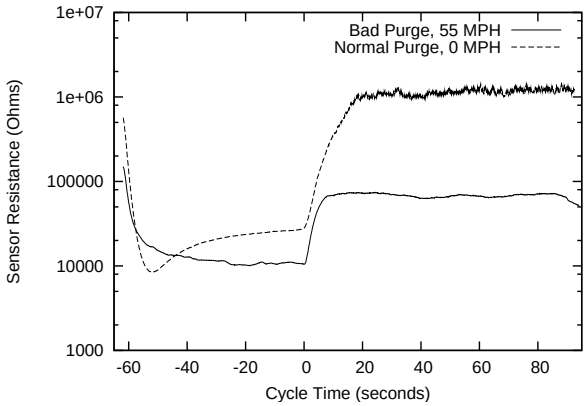

Figure 11: Exposed sensor resistance in clean (0 ppm CO) air, with 0 MPH or 55 MPH vehicle speed. When resistance does not increase over the course of a purge (cycle time < 0 seconds), the purge was not successful and the measurement cannot be used.

ating temperature during the purge cycle. Then, we explain how we address this issue by insulating the sensor and shielding it from the wind. Finally, we present measurements that verify that these steps allow accurate sensor monitoring at greater vehicular speeds.

With the Enmet RECON sensor positioned on the side of a highway, we first recorded the CO concentration while the vehicle was stopped, and then took measurements while the vehicle was moving at 55 miles per hour (MPH). The "Exposed" series in Figure 10 shows that at highway speeds, the exposed sensor differs by as much as 27 ppm from the value reported by the Enmet meter: 0 ppm.

Upon further examination of the sensor's output during the inaccurate CO measurement, we identify abnormal sensor readings in the preceding purge cycle. Figure 11 plots resistance over the course of a purge and measurement cycle while a vehicle is either stopped or moving at 55 MPH. Normally, resistance increases after an initial dip at the beginning of the purge cycle; this is because as CO is being purged from the sensor, the sensor's reading should indicate lower CO concentration with higher resistance. However, in the case of the 55 MPH output, resistance actually decreases over the course of the purge. Since the purpose of the purge cycle is to eliminate any CO bound to the sensor, this output suggests that the purge cycle is not functioning properly, causing CO to build up on the sensor across multiple measurement cycles, resulting in inaccurate measurements.

One possible explanation of the unexpected decrease in resistance during a purge at 55 MPH is that the sensor may not reach a sufficiently high temperature. This reasoning makes sense with the context in which the problem occurs; at high speeds, air flow over the sensor would remove heat from the sensor more quickly than when the vehicle were stopped.

When the vehicle is not moving and there is little or no wind, the sensor's internal heater maintains the sensor at some equilibrium temperature above the ambient temperature. Using accepted models and equations for heat loss [13], we compare heat loss in stationary and moving air conditions. We calculate heat loss to the surrounding environment using a normal convection heat loss model and Newton's Law of Cooling. Since we know the heat being dissipated by the sensor's heater (760 mW), by measuring the sensor housing's surface temperature we empirically determine a dimensionless heat flux quantity called a Nusselt Number (Nu) that represents the rate of heat loss to the environment given the physical heat transfer properties of the sensor's housing [13].

When air moves over the surface of the sensor, the rate of heat loss is determined by a different model: forced-air convection. The Nusselt Number in this model accounts for the fluid properties of the air and depends on the rate of air flow. The rate of heat loss from the sensor is proportional to $\sqrt{V}$–it increases with speed. For the condition when the node is stopped, the Nusselt Number is $Nu = 6.09$, but at at highway speeds, the Nusselt Number is 6 times larger, $Nu = 39.71$. This higher rate of heat loss prevents the sensor from reaching its optimal purging temperature, resulting in inaccurate measurements.

To counteract the chilling effect of high wind over the sensor when the vehicle is moving, we took two steps (shown in Figure 12). First, we design an insulated enclosure for the sensor with a mesh screen on the front to

Figure 12: Insulated carbon monoxide sensor assembly, with metal screen to diffuse the wind. Inset shows the internal structure of the sensor; a temperature sensor is on top of the MQ-7 carbon monoxide sensor, and the resistor that serves as a supplementary heater is attached to the bottom of the MQ-7 sensor.

reduce the rate of air flow over the sensor and increase its temperature. Second, we install a supplementary heating element underneath the sensor, which is able to further increase the temperature of the sensor. Placing the sensor in the insulated enclosure increases the equilibrium temperature of the sensor during the purge cycle by approximately 4° C, and the supplementary heater raises the temperature by an additional 5° C. When the sensor is in the insulated enclosure and the supplementary heater is activated during purge phases, we observe equivalent CO concentration readings when stopped and at 55 MPH in clean air, as shown in the "Insulated" series in Figure 10.

Our approach for insulating the sensor and providing extra heat during the purge phase enables us to collect consistently accurate CO concentration measurements under a wider range of environmental conditions, compared to the standard, exposed sensor. For scenarios where the sensor fails to purge properly despite the increased temperature in the enclosure, we filter out these measurements based on whether the purge was good: empirically determined to be an increase in sensor resistance by at least 50% over the course of the purge cycle.

5 *C3R* Data Aggregation and Distribution

A network of mobile sensors can collectively capture finer-grained patterns in actual urban pollution levels in comparison to systems with stationary nodes. The high mobility of vehicular nodes results in two effects. First, nodes are able to sample from a dynamic set of locations; having multiple vantage points reveals spatial patterns in addition to temporal patterns that can be observed with stationary nodes. Second, the nodes carry measurements to other regions of the city as they drive, which results in all nodes having larger-scoped map of pollution.

We motivate sensing from multiple vantage points using vehicular nodes by showing how pollution concentrations can vary significantly in a city over short distances and time intervals. The variations that we observe–as much as 5 ppm over distances of less than a 100 m block–would not be captured by a traditional sparse deployment of stationary monitoring stations that sample as rarely as once per hour. Then, we evaluate the possibility of a participatory, VANET-based approach for building detailed maps of pollution concentration. We adopt realistic vehicular mobility [14] and signal propagation [28] models for our simulation of a well-understood gossiping communication protocol [31]. Our results indicate that even at low vehicle penetration ratios, nodes are able to build a map at much higher temporal resolution than a typical stationary sensor, which may sample only once every hour.

5.1 Advantage of Multiple Vantage Points

Stationary sensor deployments with small numbers of nodes are unable to capture actual pollution conditions, which are characterized by significant CO concentration gradients of as much as 5 ppm over distances of only one city block (˜100 meters). From the perspective of a moving vehicle, a high-resolution dynamic map of air pollution measurements is able to reveal localized peaks of pollution concentration.

Using vehicles instrumented with our sensor apparatus, we collect two hours of air pollution measurements in downtown Chicago during the morning rush hour. The vehicles drove in a spiral pattern out from the center of the measurement area in order to obtain an approximately even spatial distribution of data points.

Figure 13 shows the diversity of CO pollution concentrations that occur within a 5-block radius of the center of downtown Chicago. Our measurements had an average CO concentration of 8.5 ppm, and standard deviation of 6 ppm. Considering the relationship of distance between observations to difference in the CO concentration (Figure 14), measurements differ by about 5 ppm at distances less than 1 km. As greater distances are considered, the difference between these measurements increases to over 8 ppm for measurements between 2 and 4 km apart. A stationary CO monitoring station, such as the one in downtown Chicago (marked in the lower left quadrant of the map), cannot see these variations. Since stationary measurement stations are typically many kilometers apart [8], they would be unable to predict pollution concentrations with error lower than 8 ppm for any point greater than

Figure 13: Map of measured CO concentration (ppm) in downtown Chicago. Each value is centered on the position at which the measurement was taken. All data points were collected within two hours. The icon of a building in the lower left corner represents the only stationary monitoring node in the Illinois air monitoring network in downtown Chicago; the next closest station is 10 km away.

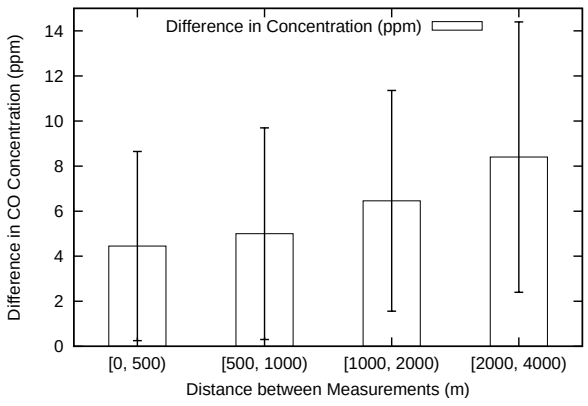

Figure 14: Difference in measured CO concentration over varying distance intervals. Over distances less than 500 meters, concentration varies by as much as 4 ppm; over longer distances (2000 to 4000 meters), concentration differs by 8 ppm.

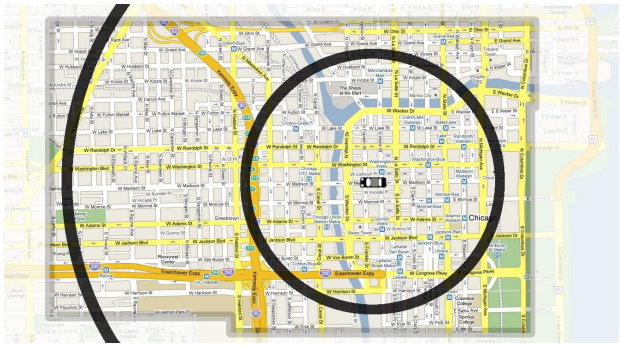

Figure 15: Simulation area of downtown Chicago. If a node is located at the center of the circle, then the inner circle indicates the mean distance of pollution measurements that are less than 15 minutes old, and the outer circle depicts the maximum range of any of those measurements.

2 km away from a station. Because 98% of Chicago is outside the 2 km radius from the downtown CO monitoring station, the stationary deployment's CO measurement error can be as high as 8 ppm for the vast majority of the city.

5.2 Data Dissemination Simulation

Via simulation of a simple gossiping protocol [31] over realistic mobility and communication models, we will show that a VANET-based system is able to provide a map of air pollution. Each node's map includes timely measurements from the surrounding area, even with a low penetration ratio of instrumented vehicles. After describing our gossiping implementation and simulator configuration, we use a latency metric to demonstrate the rapid spread of individual measurements through the network. Then, we evaluate the completeness and spatial scope of each node's pollution map.

Since the data for individual measurements are small, available bandwidth is not a limiting factor in this system. Instead, the frequency of interactions with other nodes determines measurement dissemination in the system. Empirical measurements have shown that two passing vehicles on a city street can communicate up to 2 MB of data [28]. Given that each measurement requires only 28 bytes (e.g. 8-byte double/long values for latitude, longitude, time; 4-byte single precision float for pollution measurement), two nodes could exchange 71000 pollution measurements. That many measurements would be enough for a 100-meter-resolution map of the entire City of Chicago–588 km^2 of land.

5.2.1 Configuration

We use the JiST/SWANS simulator with the STRAW vehicular mobility model [14] and empirically-determined urban signal propagation parameters [28] to achieve a realistic VANET environment. On an 8 km^2 road map from downtown Chicago (shown in Figure 15, with total road length of 266 km), we instrument 100 nodes, yielding a node density of 13 nodes per km^2. Each simulation lasts for 5 hours.

5.2.2 Gossiping Implementation

Gossiping is a well-understood approach for dissemination of data throughout a network of vehicles [31]. Since it is likely that there will be low node densities, vehicles will be disconnected most of the time. As a result, the spread of measurements through the network will rely on periodic interactions with other vehicles passing each other on the street. In our implementation, nodes periodically send a beacon message to advertise their presence to other nodes. When a node (A) receives a beacon, it probabilistically determines whether it should respond based on the number of other nodes' beacons it has received recently– this ensures that the system is scalable in terms of the node density. If it decides to continue with the interaction, it sends a digest message of the contents of its buffer to the beaconing node (B). Finally, B returns a message containing all the measurements in its buffer that A does not yet have.

Since the goal of the data dissemination system is to share recent measurements in order to maintain an up-to-date map of air pollution at every node, we remove measurements from the gossiping protocol's buffer that are older than some threshold. This limit on the temporal scope of measurements effectively bounds the spatial area from which measurements will be received, because, even if a node were to drive at highway speeds, there would be an upper bound to the distance from that node's most distal measurements based on the speed at which it moved. However, this ensures scalability of the system by implicitly limiting the scope to measurements that are relevant to the driver. Each node is equipped with a reasonably large buffer, since storage is inexpensive. In a deployment of this system, the older measurements would be stored in long-term memory for offline data processing. We adopt a threshold of 30 minutes because this is approximately the maximum observed latency for a measurement to be received at 100% of the nodes in the system. In addition, this threshold also is the warmup time for the system, after which the number of measurements in the gossiping buffer reaches equilibrium.

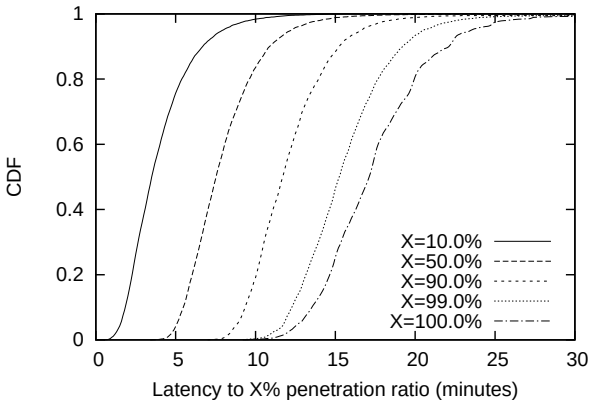

Figure 16: CDF of latency for a measurement to reach X% of the nodes in the system.

5.2.3 Evaluation: Dissemination Latency

We use latency as a metric to evaluate the dissemination of data through a disconnected network such as a VANET. In many systems (e.g. end-to-end routing protocols), there is the concept of a "recipient" node–the intended destination–for each message. In this context of disseminating air pollution measurements, however, the objective is to spread each measurement to the maximum number of nodes in the system so that each node can aggregate a dense map of measurements. We define the latency metric as the time it takes for the measurement to spread to some percentage of all nodes in the system.

Figure 16 shows that on average, every message is received by all nodes in the system in 17.3 minutes; 99.3% of the messages were received by all nodes, taking at most 30 minutes. This data highlights the explosive spread of measurements through the network; after the first 4 minutes, only 10% of the nodes have received a particular measurement. However, after 8 minutes, 50% of nodes have it, and after 12 minutes it has spread to 90% of the nodes.

5.2.4 Evaluation: Pollution Map Completeness

We evaluate the spread of measurements through the network from the perspective of individual nodes by measuring the completeness of each node's pollution map. Our metric is the proportion of recent measurements that have been received by a node, out of the total number of recent measurements in the system.

Using the specifications of our CO sensor, each instrumented node "measures" air pollution once every 2.5 minutes and injects that message into the gossiping protocol. Given the number of nodes (n) on the map, we can compute the expected number of measurements (M) that were collected in the last x minutes: $M = n * x/2.5$. Fig-

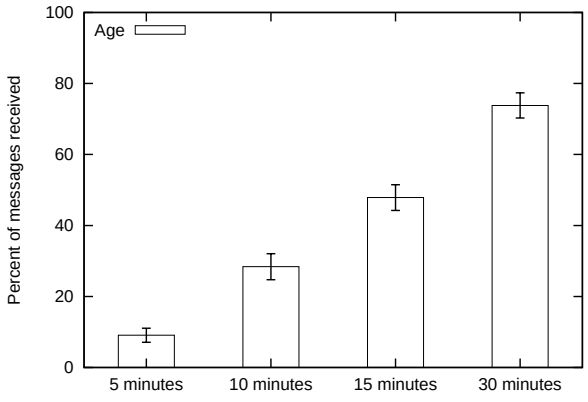

Figure 17: Mean and standard deviation of percentage of measurements received by a node out of all measurements taken by all nodes in the last x minutes.

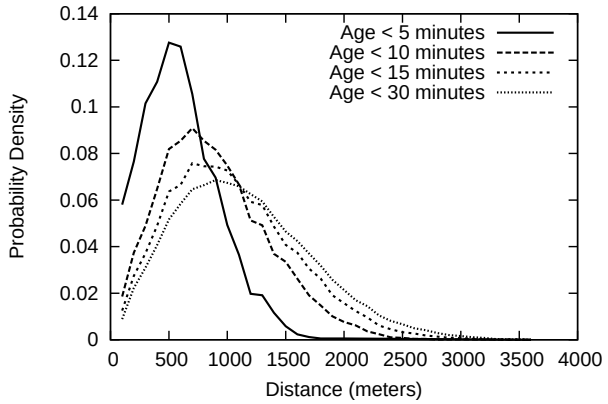

Figure 18: Probability density plot of measurement distance from a node, considering measurements that are <5, 10, 15, or 30 minutes old.

ure 17 shows the percentage of measurements received by a node, when we define "recent" measurements to be less than 5, 10, 15, or 30 minutes old.

The rapid spread of measurements means that each node has a significant proportion of the available data within as little as 15 minutes. If we only consider recent data that are less than 5 minutes old, nodes have 9% map completeness (of the available measurements). However, considering measurements less than 15 minutes old, a node has almost 50% map completeness. If we include measurements that are up to 30 minutes old, a node has nearly three quarters of the available measurements in its map.

5.2.5 Evaluation: Pollution Map Scope

There is a trade-off between the "freshness" of data and the coverage area of a map that has a measurement age threshold. Since the dissemination of measurements through the network is limited primarily by vehicular mobility and the resulting communication exchanges between passing vehicles, more recent measurements received by the node tend to come from positions closer to the node's location.

Figure 18 shows a probability density plot over measurement distance from the node, for varying measurement age limits. For only measurements taken in the previous 5 minutes, the average distance from a node's current position to a measurement location is 550 meters. This distance increases to 1100 meters when considering measurements that are up to 30 minutes old. Similarly, the distance from a node to its most distant measurement ("maximum distance") follows a similar trend. With data less than 5 minutes old, maximum distance is 1070 m. Considering measurements up to 10 minutes old, maximum distance increases to 1935 m, and with 30 minutes of measurements, the maximum distance is 2572 m.

To illustrate the physical scale of these values, Figure 15 shows the radii of the average (990 m) and maximum (2369 m) distances consider measurements taken in the last 15 minutes, centered around the black dot that represents a vehicular node. All the nodes' measurements would be within the larger outer circle (about a 20-block radius), while approximately half of its measurements would be within the smaller circle (an 8-block radius).

6 Conclusion

This paper presented *C3R*, a vehicular-based, participatory approach to urban pollution monitoring. The work is motivated by the need for high-resolution, up-to-date information to support policy makers and public awareness. We described the *C3R* approach and its current realization. *C3R* nodes are based on low-cost Carbon Monoxide (CO) sensors. We showed how these low-cost sensors can be adapted for operation in the challenging vehicular context. We have used *C3R* to survey the distribution of CO in a region, showing that any practical approach based on stationary nodes cannot capture the variability observed in urban environments. Finally, we evaluated the feasibility of the basic *C3R* gossiping-based model for data aggregation and distribution, focusing on the freshness and coverage of the data used to generate each node's map of pollution concentration. An interesting and unique advantage of a participatory approach to pollution monitoring is that it makes possible a form of "soft" control through a direct feedback channel of every driver's contribution. Finally, while our work has so far concentrated on monitoring concentration levels of CO, we believe our results are easily generalizable to most other pollutants (from both natural and anthropogenic sources). We have started to explore

several extensions to the basic *C3R* platform described.

References

[1] AIRNow. http://airnow.gov/.

[2] Conceptual strategy for ambient air monitoring. http://www.epa.gov/ttn/amtic/files/ambient/criteria/reldocs/799stegy.pdf.

[3] Cook county environmental control ordinance. http://cookcountygov.com/Agencies/cc_envcont_ord.pdf.

[4] Enmet corporation. http://www.enmet.com/.

[5] GarbageWatch. http://garbagewatch.com/.

[6] Grimm technologies, inc. environmental dust monitor model 1.180. http://www.dustmonitor.com/Occupational/180.htm.

[7] Hanwei Electronics MQ-7 CO Sensor. www.hwsensor.com/English/PDF/sensor/MQ-7.pdf.

[8] Illinois ambient air monitoring network plan for 2009. http://www.epa.state.il.us/air/monitoring/network-plan-2009.pdf.

[9] *Air quality guidelines. Global update 2005.* World Health Organization, 2006.

[10] *Preventing disease through healthy environments: Towards an estimate of the environmental burden of disease*. World Health Organization, 2006.

[11] Illinois annual air quality report, 2007. http://www.epa.state.il.us/air/air-quality-report/2007/index.html.

[12] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. Srivastava. Participatory sensing. In *Proc. of WSW*, 2006.

[13] Y. A. Cengel. *Heat Transfer – A Practical Approach*. WCB/McGraw-Hill, 1998.

[14] D. R. Choffnes and F. E. Bustamante. An integrated mobility and traffic model for vehicular wireless networks. In *Proc. of ACM VANET*, 2005.

[15] J. Chou. *Hazardous Gas Monitors: A Practical Guide to Selection, Operation and Applications*. SciTech Publishing, 2000.

[16] B. Croxford, A. Penn, and B. Hillier. Spatial distribution of urban pollution: civilizing urban traffic. *Science of The Total Environment*, 1996.

[17] S. Eisenman, E. Miluzzo, N. Lane, R. Peterson, G.-S. Ahn, and A. Campbell. The BikeNet mobile sensing system for cyclist experience mapping. In *Proc. of ACM SenSys*, Sydney, Australia, November 2007.

[18] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan. The Pothole Patrol: using a mobile sensor network for road surface monitoring. In *Proc. of ACM/USENIX MobiSys*, Breckenridge, U.S.A., June 2008.

[19] M. Ghanem, Y. Guo, J. Hassard, M. Osmond, and M. Richards. Sensor grids for air pollution monitoring. In *3rd UK e-Science All Hands Meeting*, 2004.

[20] B. Hoh, M. Gruteser, R. Herring, J. Ban, D. Work, J.-C. Herrera, A. Bayen, M. Annavaram, and Q. Jacobson. Virtual trip lines for distributed privacy-preserving traffic monitoring. In *Proc. of ACM/USENIX MobiSys*, Breckenridge, CO, June 2008.

[21] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: a distributed mobile sensor computing system. In *Proc. of ACM SenSys*, Boulder, Colorado, USA, 2006.

[22] N. A. H. Janssen, B. Brunekreef, P. van Vliet, F. Aarts, K. Meliefste, H. Harssema, and P. Fischer. The relationship between air pollution from heavy traffic and allergic sensitization, bronchial hyperresponsiveness, and respiratory symptoms in dutch schoolchildren. *Environ Health Perspect*, (12), 2003.

[23] U. Lee, E. Magistretti, B. Zhou, M. Gerla, P. Bellavist, and A. Corradi. MobEyes: smart mobs for urban monitoring with vehicular sensor networks. *IEEE Wireless Communications*, 13(5), October 2006.

[24] T. S. Lena, V. Ochieng, M. Carter, J. Holguin-Veras, and P. L. Kinney. Elemental carbon and pm2.5 levels in an urban community heavily impacted by truck traffic. *Environmental Health Perspectives*, (10), 2002.

[25] P. Mohan, V. N. Padmanabhan, and R. Ramjee. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. In *Proc. of ACM SenSys*, 2008.

[26] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. Trafficview: traffic data dissemination using car-to-car communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(3), 2004.

[27] R. North, M. Richards, J. Cohen, N. Hoose, J. Hassard, and J. Polak. A mobile environmental sensing system to manage transportation and urban air quality. In *ISCAS*, 2008.

[28] J. S. Otto, F. E. Bustamante, and R. A. Berry. Down the block and around the corner – the impact of radio propagation on inter-vehicle wireless communication. In *Proc. of ICDCS*, 2009.

[29] S. Reddy, K. Shilton, J. Burke, D. Estrin, M. Hansen, and M. Srivastava. Evaluating participation and performance in participatory sensing. In *UrbanSense*, 2008.

[30] W. Tsujita, H. Ishida, and T. Moriizumi. Dynamic gas sensor network for air pollution monitoring and its auto-calibration. In *IEEE Sensors*, 2004.

[31] A. Vahdat and D. Becker. Epidemic routing for partially connected ad-hoc networks. Technical report, Duke University, 2000.

[32] L. E. Venegas and N. A. Mazzeo. Carbon monoxide concentration in a street canyon of buenos aires city (argentina). *Environmental Monitoring and Assessment*, November 2000.

[33] J. Yoon, B. Noble, and M. Liu. Surface street traffic estimation. In *Proc. of ACM/USENIX MobiSys*, San Juan, Puerto Rico, June 2007.