

Eclipse: An End-to-End Platform for Low-Cost, Hyperlocal Environmental Sensing in Cities

Madeleine I. G. Daepf
Microsoft Research
Redmond, Washington
USA
mdaepf@microsoft.com

Alex Cabral
Harvard SEAS
Cambridge, Massachusetts
USA
acabral@g.harvard.edu

Vaishnavi
Ranganathan
Microsoft Research
Redmond, WA, USA

Vikram Iyer
University of Washington
Seattle, Washington, USA

Scott Counts
Paul Johns
Asta Roseway
Microsoft Research
Redmond, Washington
USA

Charlie Catlett
University of Illinois
Discovery Partners
Institute
Chicago, Illinois, USA

Gavin Jancke
Darren Gehring
Chuck Needham
Curtis von Veh
Microsoft Research
Redmond, Washington
USA

Tracy Tran
Lex Story
Gabriele D'Amone
Bichlien H Nguyen
Microsoft Research
Redmond, Washington
USA




Figure 1: The Eclipse platform includes a modular environmental sensing device (left panel) that can augment existing urban infrastructure such as bus shelters (middle panel). The real-time web interface shows air quality data at over 100 locations in Chicago, Illinois, USA on July 20, 2021, during a smoke event from western wildfires (right panel).

ABSTRACT

This paper presents Eclipse, a platform for low-cost urban environmental sensing using solar-powered and cellular-connected devices. Dense sensor networks promise to monitor pollution at fine spatial and temporal resolutions, yet few cities have actually implemented such networks due to high costs and limited accuracy. We address these barriers by developing an end-to-end framework for urban air quality sensing with minimal infrastructure requirements. We designed an unobtrusive device that collects data on fine particulate matter (PM_{2.5}), temperature, relative humidity, and barometric pressure. A modular design further includes four low-cost gas sensors — Ozone (O₃), Nitrogen Dioxide (NO₂), Sulfur Dioxide (SO₂), and Carbon Monoxide (CO) — selected based on local priorities. We deployed 115 devices across Chicago, reliably collecting data for over 90% of expected sensor-hours from July 2 - September 30, 2021.

We further developed a calibration strategy that reduced errors by 41.2 - 98.8%, improving accuracy to levels recommended for hotspot detection (PM_{2.5} and O₃) or education (NO₂ and SO₂). Through this work, we offer insights on the real-world deployment of a replicable, large-scale, end-to-end platform for hyperlocal urban environmental sensing.

KEYWORDS

Internet of Things, Air Pollution, Smart Cities, Sensor Networks

1 INTRODUCTION

Over 4 billion people — more than 55% of the world's population — live in cities, and this number is expected to grow to nearly 70% by the year 2050 [58]. Although rapid urbanization has economic and

other social benefits, it has also exposed more people to environmental hazards including air pollution – the largest environmental contributor to mortality [60]. Poor air quality is linked to a number of adverse health effects, including heart and lung disease, as well as asthma [19, 36, 61]. To monitor environmental pollutants, regulators and policymakers rely on data from regulatory equipment managed by government agencies and research institutes. However, highly accurate regulatory monitors are expensive, large, and require special expertise for maintenance. As a result, regulatory networks are geographically sparse and thus unable to capture known variability that occurs at finer spatial resolution [30, 48]. For urban public health and planning applications – which require an understanding of intra-urban spatial inequities and evaluations of policies over time – there is a strong need to collect real-time data at finer spatial resolutions.

To address this need, we present Eclipse: an end-to-end platform for low-cost, hyperlocal environmental sensing in cities. Key components of the Eclipse platform include: a low-power, cloud-connected, solar-powered, multi-pollutant sensing device; a deployment and maintenance strategy; and a site selection process. An Azure cloud back-end includes a complete data processing pipeline and application programming interface (API). Our calibration strategy takes noisy raw, low-cost sensor data and makes it usable for analytic applications. Finally, we also create a public-facing website to give users real-time data from an ongoing 115-node deployment in Chicago, USA. Because our physical device and network connectivity are built with off-the-shelf solutions, the core contribution of this work is showing how these elements can be used to design a scalable environmental sensing solution for long-term, real-world deployments that address the needs of cities. We thus present findings and lessons learned from a months-long deployment in a major city that incorporates physical sensors, cloud analytics, and a real-time interface for the public.

Designing a citywide environmental sensing system requires addressing numerous challenges. Despite a proliferation of demonstration projects promising an “internet of environmental things”, few cities have successfully deployed dense sensing networks due to four key barriers:

- (1) Deploying a large network poses substantial costs both for hardware and for ongoing maintenance [11].
- (2) Many existing devices require dedicated infrastructure – hard-wiring for power and connectivity – that severely constrains the set of possible deployment locations and dramatically increases installation and maintenance costs [1].
- (3) There is a trade-off between affordability and accuracy: low-cost sensors are subject to errors due to interference and drift, among other factors [40].
- (4) Deployments occur in shared public spaces, raising a critical need for buy-in from city and community leaders [24], including mechanisms for participation and accessible data [18].

The Eclipse platform is designed to address these barriers by leveraging recent advances in low-power sensing and calibration to support the deployment of a dense, real-time urban environmental sensing network designed in collaboration with local stakeholders. We report on six important components of the work. First, we developed and optimized the Eclipse hardware for ultra-low power

operation including the use of harvested solar power. Second, nodes communicate directly to the cloud via existing cellular networks, which eliminates the need for network setup and allows for remote monitoring and software updates. As a result, the Eclipse devices can be deployed at a multitude of locations in a city by augmenting existing street furniture, light poles, or other physical infrastructure. This approach increases the number of possible sites and significantly reduces maintenance costs, allowing devices to be easily moved and replaced as part of existing upkeep procedures. Third, we created a complete Azure cloud backend and data processing pipeline. Data is taken into a SQL database and Stream Analytics and exposed to a high-level API that enables users to access data as well as Power BI analytics. Fourth, we developed a calibration strategy that reduces errors by 41.2 - 98.8% through adjustments for relative humidity, temperature, and the complete array of pollutants measured. Fifth, we developed a user-friendly website to visualize data from Eclipse devices in real time. The site is accessible via a direct URL or QR codes posted at augmented bus shelters and shows the recently recorded values from Eclipse devices. Finally, our work was the product of multi-sectoral collaboration: we worked with JCDecaux Chicago, the local affiliate of JCDecaux SA – the world’s largest outdoor advertising company, which installs and operates bus shelters and other streetscape structures in cities globally – to install and monitor Eclipse devices on bus shelters throughout the city. We also collaborated with the City of Chicago, the academic Array of Things initiative, and a network of local environmental justice organizations supported via the Environmental Law and Policy Center. These partners’ needs informed key design decisions including the focus on low-cost hardware with minimal infrastructure needs and the development of easily accessible data visualization tools, ensuring that the Eclipse approach reflects the real-world priorities and constraints of diverse stakeholders.

We deployed Eclipse in three overlapping phases: a pilot deployment of 6 devices beginning in April 2021, a regulatory co-location deployment of 9 devices beginning in May 2021, and a full deployment of 100 additional sensors beginning in July 2021; this paper shares findings from the first three months of the network’s operation (July 1 until September 30, 2021). From these deployments, we report results on platform reliability, data quality, and website engagement. Notably, we collected readings for >90% of expected sensor hours, minimizing data loss despite initial start-up issues that led to the relocation of $\approx 10\%$ of devices. Moreover, calibration improved the accuracy of the raw data to levels consistent with U. S. Environmental Protection Agency (EPA) guidelines for hotspot characterization ($PM_{2.5}$) or monitoring for education and awareness (O_3 and NO_2) [62]. We saw limited traffic to the website via QR codes, suggesting a need to reduce barriers between residents and sensor data; nevertheless, interviews and ongoing engagements with city and community partners offer evidence of strong local interest in and support for the monitoring network.

The key contribution of our work lies in identifying challenges associated with long-term real-world deployments in cities and in developing a deployment template to address those challenges. Although the focus of this paper is on air quality sensing, the Eclipse approach offers a replicable framework for the deployment of urban environmental sensing networks more generally. We describe the development of low-cost hardware with minimal infrastructure

Network	No. of Sensors	Timeline	Data Transmission	Power Source	Public-Facing, Real-Time Data
RAMP [54]	40+	1 year	Cellular	Rechargeable battery	No
AirU [34]	80+	2016 –	WiFi	5V battery pack	No
BEACO ₂ N [52]	50	2016 –	Ethernet, WiFi, Cellular	120V Outlet	Yes
Array of Things [11]	150	2015-2020	Cellular	120V Outlet	Yes
AirBox [13]	1000+	2016 –	WiFi	12V Outlet	Yes
Breathe London [7]	100+	2018 –	Cellular	Solar Power	Yes
Eclipse	115	July 2021 –	Cellular	Solar Power	Yes

Table 1: A comparison of Eclipse to prior and existing large-scale (40+ sensors), long-term (several months of data), “full-stack” stationary networks including hardware, network design, and data access for urban air quality. Although Breathe London uses commercially available sensors (every other project developed their own), it is included as it is the most similar to Eclipse.

needs and show how our approach can easily be incorporated into existing urban upkeep processes, enabling rapid deployment and reducing maintenance costs. Over the course of three months, we find that our framework addressed known problems around power and connectivity; however, new barriers emerged regarding city-wide network availability, communicating findings to the public, and harvesting solar energy in built-up areas. Nevertheless, our experience shows that low-cost, citywide sensing networks are feasible and uncovers key lessons and opportunities for environmental sensing at the urban scale.

2 RELATED WORK

We review prior work on low-cost sensor networks in urban areas that are comparable to our deployment in Chicago. These networks comprise mobile sensors, stationary sensors, or a combination of the two. We evaluate these networks based on the needs of cities, which include scalability, minimal infrastructure requirements, and relevance for core planning and public health applications.

Mobile sensor networks. Mobile sensor networks involve the deployment of low- or medium-cost sensors on existing city transit and vehicle fleets [2, 41], vehicles and pedestrians who volunteer via crowdsourcing [14, 32], or purposefully outfitted monitoring vehicles [3, 27, 41]. These approaches are promising and have produced interesting early findings [12], but they are also subject to important limitations. First, because many mobile monitoring networks rely on vehicles [2, 3, 20, 27], sensors are deployed in the middle of streets. Although this is an important site for characterizing near-source emissions, it is not where pedestrians breathe and thus are exposed to emissions. As a result, such mobile data is less useful for public health applications such as cumulative exposure studies or health impact assessments. Mid-street measurements also cannot capture the effect of planning interventions that disrupt the flow of emissions from streets to surrounding areas (e.g. barriers or vegetative buffers), limiting the extent to which they can be used to evaluate or guide new city policies. Finally, because they lack extended and continuous temporal measurements, mobile fleets can miss important short-term events and may not allow the assessment of spatial inequities in emission patterns during pollution events. Mobile networks, although promising for research applications, are thus limited in their utility for the routine monitoring applications needed to inform urban public health and planning practice.

Quasi-Regulatory Networks. Cities seeking to monitor routine environmental exposures and inequities have tended to turn towards dense, stationary networks [7, 44]. The New York City Community Air Survey, for example, deploys temporary (2-week) networks of expensive research-grade sensors 4 times per year [44], a strategy that has been highly effective in supporting the implementation and evaluation of new pollution mitigation policies [35]. However, data are stored locally and analyzed manually, so the network does not offer insight on real-time events and events that occur outside of the limited seasonal deployment.

Crowd-sourced Networks. Other projects rely on crowd-sourced Internet of Things (IoT) networks: Taiwan and 29 countries saw a proliferation of over 2500 sensors via a participatory framework using various devices connected to one common platform [13], and there has been a proliferation of research using data from increasingly common PurpleAir nodes [6, 42, 47]. Because these sensor networks require hardwired power and WiFi – adding significant cost for city governments to deploy them at scale – they rely on sensors installed by private citizens who volunteer their resources. However, this approach introduces systematic bias: PurpleAir nodes, for example, are more likely to be located in socioeconomically advantaged versus in disadvantaged neighborhoods [21]. Crowd-sourced networks such as these have thus been poorly suited for evaluations of economic or racial inequities in exposures, which is problematic because health equity is becoming a key monitoring priority for many urban public health departments [37].

Academic Networks. Academic initiatives tend to be smaller or short in duration [45, 46], offering limited insight into the challenges of deploying at scale for an extended period across an entire city. However, recent advances include a 40-node network in Pittsburgh [54], an 80-node network in and around Salt Lake City [5, 34], and the 50-node BEACO₂N network in San Francisco Bay [52, 57], all of which required hardwiring for power or depended on WiFi connectivity, features that significantly increase maintenance costs or limit deployment locations. Moreover, these networks generally rely on private volunteers or researchers to address routine maintenance needs.

The 150-node Array of Things initiative in Chicago did incorporate collaborations with city and community stakeholders for network design and maintenance [24, 55]. However, air quality was just one of many measurements [10] that the Array of Things

evaluated and the particulate matter sensors were susceptible to hardware failure (optics misalignment) during transport and installation. These problems were confirmed after many nodes had been deployed, when replacing the node would have interfered with data collection by the many working sensors for other parameters. The Array of Things was also constrained by the need to hardwire devices for power: the most ubiquitous source of power in Chicago is at intersections where signaling and control equipment are located—but these are also locations where air pollution levels can be skewed by road traffic. By contrast, energy-harvesting devices can be deployed in diverse locations including within green spaces, on waterfronts, or in residential areas without major streets. The challenges encountered by the Array of Things thus highlighted a need for modular, low-cost devices that would be both easy to install and inexpensive to replace when necessary – a constraint that became a driving design consideration for Eclipse.

The most similar initiative to Eclipse is Imperial College London’s *Breathe London* initiative [49]. A network of 100 sensor nodes in London, *Breathe London* has been running since 2018. The network uses commercially-produced Clarity nodes to report calibrated $PM_{2.5}$ and O_3 data via an interactive interface (though the network was originally deployed with commercial AQMesh Nodes, replaced after 2020). Similar to the Eclipse device, the Clarity nodes are solar-powered and LTE-connected; however, the Clarity node is larger and heavier than the Eclipse device – limiting the set of possible installation locations – and collects data on only two pollutants ($PM_{2.5}$ and NO_2). Moreover, although *Breathe London* has benefited epidemiological research as well as city policy [7], the need for maintenance and monitoring has limited replication in other cities, highlighting the need for a simple end-to-end deployment framework.

Eclipse aims to fill several important gaps relative to prior work. First, the Eclipse devices are designed using a combination of technologies that requires no wiring for power and that can use cellular networks rather than requiring the setup of new local area networks. These characteristics address city stakeholders’ needs for flexible deployments with minimal infrastructure requirements. Second, we address the need for equitable coverage through both the reduced setup complexity and through a network design with both citywide coverage and additional sites in environmental justice neighborhoods. Finally, the low-cost platform, inclusion in existing urban maintenance routines, collaborative network design, and accessible data dashboard constitute a deployment framework that could easily be replicated in other cities.

3 THE ECLIPSE SENSING ARCHITECTURE

This section presents the key design considerations and implementation details involved in developing the low-cost and low-overhead Eclipse device (Fig. 2). Based on many of the lessons learned from the Array of Things project and other prior work, we prioritized four requirements for the sensor module design:

- Low power operation & energy harvesting
- City-wide wireless communication
- Modular environmental sensing
- Air-flow design & weather-proofing

This section also describes the cloud-based analytics and visualization used to enable public access to real-time data streams.

3.1 Low-Power Operation & Energy Harvesting

Historically, one of the key constraints on urban sensor networks has been the limitation of deployment sites to wired mains [1]. Although cities have power available at streetlamps, traffic signals, and other similar locations, past experiences by members of our team [11] as well as in other, similar initiatives [44] have shown that relying on these access points increases costs – due to the specialized labor required both for installation and maintenance – and limits the locations at which samples can be taken. The Eclipse sensor device eliminates this key, wired-power infrastructure dependency by leveraging solar energy harvesting. We used a small (10×13 cm) high-efficiency solar panel (Voltaic Systems P126) with a UV and waterproof PTFE coating. The cell provides 6 Watts peak power and charges a Lithium polymer (LiPO) battery with a capacity of 2000 mAh as shown in the power module of Fig. 2. A BQ25570 energy-harvesting IC along with a dedicated LiPo battery charger IC make up the main energy-harvesting and power monitoring unit. We designed our hardware, communication, and sensors to optimize for low-power operation. The NRF9160 microcontroller is operated in a duty cycled mode to leverage its deep sleep mode which consumes as little as 40 μA between measurements. The device samples its four electrochemical gas sensors, which also consume microwatts of power, every 60 seconds. In contrast, the particulate matter (PM) sensor (Sensirion SPS30) consumes up to 80 mA of current while sensing as it uses an internal fan to circulate air. To optimize overall system power we sampled the PM and transmitted data off the device every 5 minutes. The average current draw of the device over a 24-hour period is 4 mA; without the solar harvesting the 2000 mAh battery can sustain the Eclipse device at this sampling rate for approximately 15 days, allowing the device to operate robustly even through prolonged periods of low light.

3.2 City-Wide Wireless Communication

Many existing low-cost environmental sensing networks use Low Power Wide Area Networks (LPWAN) such as LoRa, Sigfox, and Zigbee. These approaches introduce four main challenges associated with the need to set up a large scale network of dedicated wireless access points (with IoT radios to forward data to the Internet from sensor devices) across a metropolitan area. First, each access point requires a dedicated power and backhaul solution. Although the access point itself is relatively inexpensive, installing the wired power and Internet connectivity can incur additional setup and maintenance costs as described above. Second, setting up access points requires negotiating access to physical infrastructure such as rooftops and other locations. Third, each of these components requires regular maintenance from a dedicated team of trained professionals, increasing labor and training costs. Lastly, IoT networks are primarily designed for low-bandwidth, one-way communication, which cannot support access to the devices for remote software updates or diagnostics. We thus designed Eclipse devices to instead utilize existing commercial 4G LTE-M cellular networks. This approach eliminates the need for maintenance of

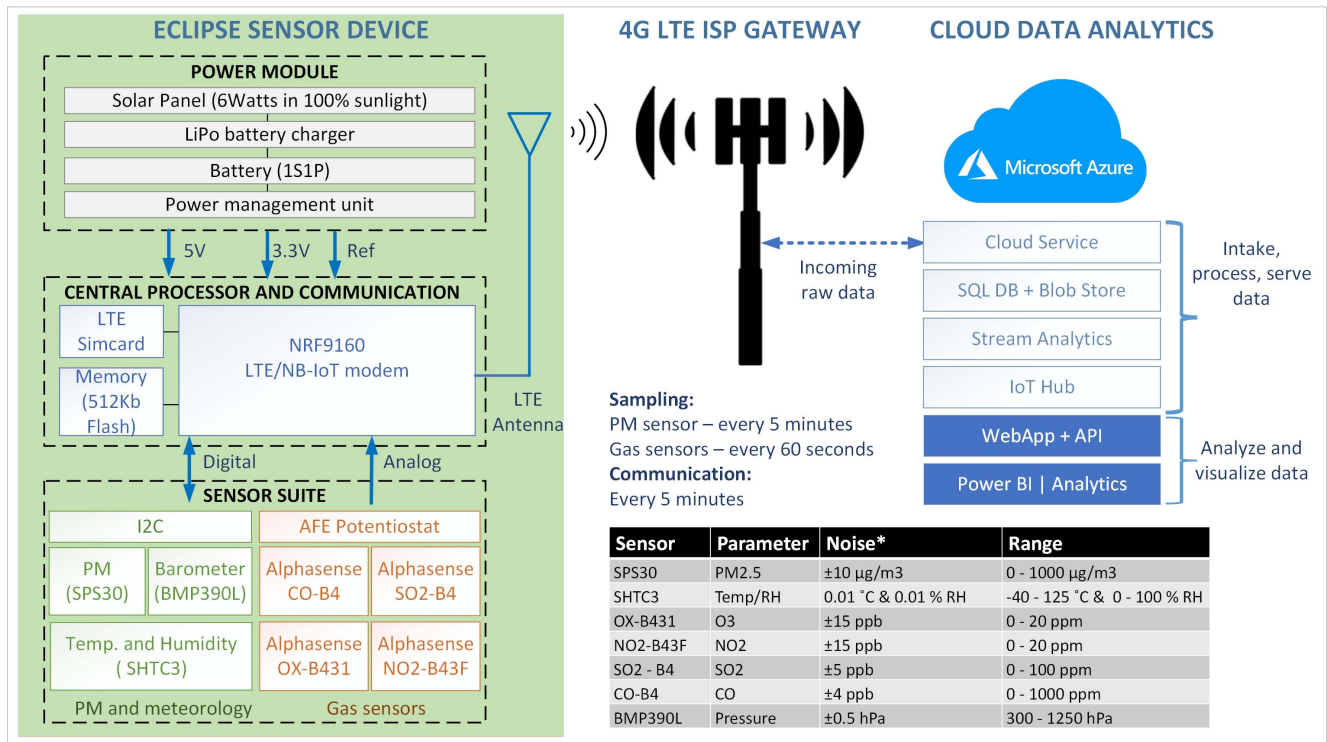


Figure 2: The Eclipse device architecture (left) and data pipeline (right). The block diagram on the left presents the three important modules in the device design: 1) the power module that supports energy harvesting, 2) the central processor and communication module that enables low-power operation with LTE communication and 3) the sensor suite to measure air quality and environmental parameters. The figure also describes data communication via 4G LTE ISPs and cloud data analytics (right). The sampled data is communicated every 5 minutes, over LTE connection, directly to Azure cloud. Real-time processing, analysis and visualization are performed using Azure services and the air quality data is made accessible via an interactive website. Finally, the table details the specifications of all sensors used in the Chicago deployment. *(Note that noise for the SPS30 is $\pm 10 \mu\text{g}/\text{m}^3$ for readings $\leq 100 \mu\text{g}/\text{m}^3$ and $\pm 10\%$ for readings $> 100 \mu\text{g}/\text{m}^3$.)

the wireless network itself and gives nodes the ability to upload data directly to the cloud.

We designed our devices using the NRF9160 “microcontroller plus modem” which can operate at sufficiently low power to run on energy harvested from a small solar panel. We also implemented over-the-air (OTA) programming, eliminating the need for additional infrastructure for software upgrades and maintenance. The device also has on-board flash memory to buffer data during times of connectivity failure. The central processor and communication block of Fig. 2 provides a high level summary of the processing and communication elements in the Eclipse sensor.

3.3 Modular Environmental Sensing

Our device includes a modular, adjustable set of environmental sensors (detailed in Fig. 2). Given our focus on air quality, we selected sensors to measure five of the six criteria air pollutants identified by the U.S. National Ambient Air Quality Standards (NAAQS) [59]. The Eclipse devices evaluate $\mu\text{g}/\text{m}^3$ of fine particulate matter (PM₁, PM_{2.5}, and PM₁₀) using a Sensirion SPS30 sensor, which outperforms other low-cost sensors in laboratory [56] and field tests [51], performing similarly to more expensive devices in field validation

studies in detecting PM₁ and PM_{2.5} albeit with poorer results for the detection of PM₁₀ [53]. For gases, we included electrochemical sensors from Alphasense (CO-B4, NO₂-B43F, and OX-B431) given evidence of their pre-calibrated accuracy and stability across a range of environmental conditions and ambient gas concentration mixes [17]. We also included the Alphasense SO₂-B4 sensors, despite prior reports of poor field performance [25, 26], because of local concerns regarding emissions in Southeast Chicago. Finally, the Eclipse device also monitored relative humidity (RH), barometric pressure, and ambient temperature.

3.4 Air-Flow & Weather-Proofing

The Eclipse device was designed with an unobtrusive leaf-shaped exterior as a reference to its environmental applications. To ensure that the shape did not compromise consistent air flow across the sensors, we designed the enclosure of the Eclipse device to optimize the airflow path across the gas sensors. We performed computational fluid dynamics (CFD) simulations of airflow to evaluate non-homogeneity, stagnation, or recirculation. We then introduced the following specific modifications, shown in Fig. 3:




Figure 3: The top enclosure half with the left and right air flow channels. The channels are designed to optimize steady non-stagnant airflow, over the gas sensors, from the bottom inlets to the top outlets.

- Overall channel volume was reduced to a minimum by adding curb to fill lateral spaces and modifying the profile of inlets and outlets to avoid dead volumes.
- A curved conduit was introduced to guide air flow from the inlet over the sensor array to the vent. This increased uniformity of the flow over the sensors and minimized the width of potential recirculating zones.
- Sensors were oriented vertically to promote upward airflow.

To ensure weather and pest resistance, we shielded the inlet and outlet vents with a fine mesh and weather-proofed the device, fully sealing the enclosure except around the sheltered airflow inlets. Separate left and right airflow channels, inside the top half of the enclosure, connect these inlets to the sheltered vents at the top on either side. A magnetic reset enabled maintenance crews to reset the device in the field without opening the enclosure. We 3D printed the device enclosure with Nylon and apply a hydrophobic and inert surface coating (Microcure DTO, Cytonix) to prevent the target gases from reacting with the enclosure material.

3.5 Cloud Data Analytics

We leverage cellular LTE communication to transmit sensor data from each device. As illustrated in Fig. 2, the raw data gets ingested by an Azure cloud pipeline, where it is first stored in a SQL database and can be managed using Stream Analytics and Azure IoT services. We further process the data for analysis using Power BI and Azure ML; the data is then presented for public access through an API and a custom website (discussed in detail in Section 5).

4 DEPLOYMENT AND CALIBRATION

This section describes the steps taken to deploy, maintain, and calibrate our network of 115 sensors across the city of Chicago. We discuss our phased deployment and tools for device monitoring.




Figure 4: Three Eclipse devices were co-located with air pollution monitoring equipment at each of three EPA sites in the city of Chicago throughout the deployment to enable the evaluation of accuracy and the development of data calibration models.

We then detail the approach used to evaluate and improve accuracy through sensor calibration. Finally, we describe our integral partnership with local agencies and organizations.

4.1 Deployment Strategy

Our deployment was carried out in three phases. First, we launched a small pilot program in April, 2021, to ensure the feasibility of the entire deployment. The pilot consisted of 6 devices deployed at 4 bus shelters selected by local partners and 2 bus shelters across the street from other sites to evaluate sensor precision. JCDecaux employees installed the devices on April 12th and 13th, 2021.

Second, to evaluate the accuracy of the Eclipse devices compared to research grade monitors, we co-located 3 devices at each of 3 EPA regulatory stations ($N = 9$ Eclipse devices total). We installed 3 devices on April 14th and 6 devices on May 10th, attaching the devices to monitoring stands (Fig. 4).

Third, we designed the main deployment using a sampling and community engagement process based on the methods of the NYC-CAS [44]. We deployed 80 devices across the city using a stratified random sampling approach based on traffic and population density. We then allocated 20 additional devices to locations selected by local environmental justice organizations (Section 4.4). Devices were installed by three teams of workers over a two-week period between June 29th and July 7th. Teams successfully installed 93 devices during this period; an additional 7 devices were installed in July and August due to delays in community groups selecting locations. The final network included 115 devices at locations shown in Fig. 5. The number of devices deployed was limited by cost and manufacturing constraints; nevertheless the final coverage ($M = 115$, 0.2 sensors/ km^2) is at the upper end of the sample sizes commonly used in studies of intraurban air quality [31] and in citywide networks (Table 1).

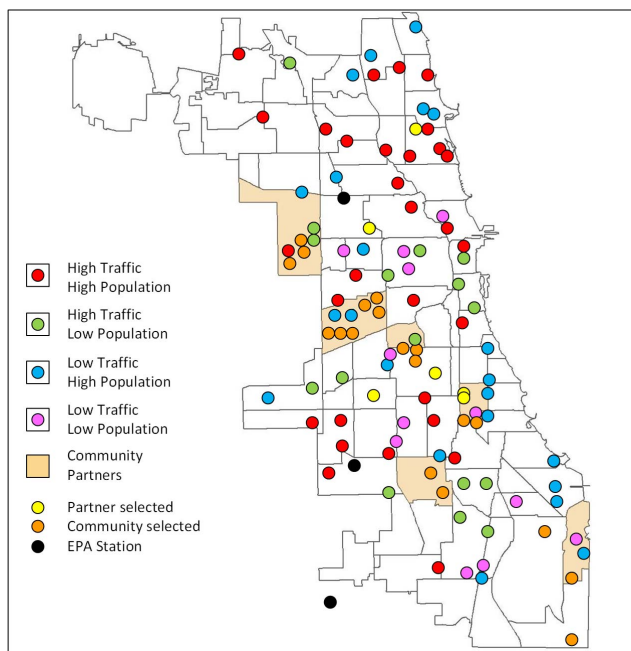


Figure 5: Site selection. We allocated 80 devices to sites via a stratified random sampling approach. An additional 20 sites were chosen by community organizations in environmental justice neighborhoods (shaded) and local partners requested 6 additional sites. Finally, black points indicate the 9 devices co-located at each of 3 regulatory monitoring stations.

4.2 Maintenance Toolkit

We created several network monitoring tools: a website showing each device’s most recent received signal strength indicator (RSSI), battery voltage, and $PM_{2.5}$; a PowerBI with scripts to visualize battery issues; and automated daily emails indicating any devices that stopped reporting in the previous 48 hours. We reached out to the local maintenance team whenever these tools identified issues.

4.3 Sensor Evaluation and Calibration

Strategies for evaluating and improving accuracy are a necessary component of any environmental sensing framework because concerns about data quality have played a major role in impeding the widespread use of low-cost air quality sensors [39, 43]. In particular, field validation studies show large errors due to interference from meteorological parameters or other pollutants and to drift in sensor readings over time [9, 26, 43, 53]. This section describes our approach to evaluate and improve sensor accuracy as part of the Eclipse framework; results are discussed in Section 6.2 and Fig. 9.

4.3.1 Regulatory Co-Locations. The Eclipse network includes 3 devices co-located at each of 3 EPA regulatory stations ($N = 9$ devices total). All three stations collect hourly data on $PM_{2.5}$ from a Beta Attenuation Monitor (BAM); two stations additionally collect hourly O_3 data from Teledyne monitors (method 087); and one station reports hourly NO_2 and SO_2 data, collected from a Thermo-Scientific Monitor (method 074) and a Teledyne Monitor (method 100), respectively [33]. Although $PM_{2.5}$ and O_3 data are available

in real time to EPA Air Now,¹ NO_2 and SO_2 data are only available via the EPA Air Quality System (AQS) with an approximate three-month lag for Quality Assurance/Quality Control (QA/QC) processes.² There is no regulatory monitor for CO in Chicago, so this study did not evaluate the accuracy of the CO sensor.

4.3.2 Evaluation of Raw Data. We evaluated the precision and accuracy of the raw data by calculating the Coefficient of Variation (CV) and Normalized Root Mean Squared Error (NRMSE) following EPA guidelines [62]. We assessed linearity with R^2 as well as via the intercept and slope from a simple linear regression [22, 23, 62].

4.3.3 Development of Calibration Models. We developed calibration models to adjust hourly sensor readings for potential sources of interference and, for gas sensors, drift. We accounted for temporal dependencies in the data, grouping data by day before allocating the dataset into training (70%) and test sets (30 %). We used 5-fold cross-validation, with data grouped by day, to implement and evaluate three models that have been shown to perform well for the calibration of low-cost sensors: multivariate linear regression, random forests, and gradient boosting [15, 43]. Following previous research [8, 38, 63, 64], we evaluated models controlling for relative humidity (RH) and temperature (T); we further assessed whether using exponential transformations, b-splines, or interactions with other variables improved the models’ fit. To address concerns related to interference by other pollutants, for each outcome, we included raw measures of the other four pollutants measured as model inputs. Finally, to account for bias and sensor drift, the calibration models for the gases also included a zero value offset: a variable equal to the raw eclipse reading at the most recent time for which EPA stations reported levels at or close to zero. These offsets were updated over time as new readings were collected. Our $PM_{2.5}$ models were additionally fitted using the mass concentration of particles in each size bin ($PM_{0:1}$, $PM_{1:2.5}$, $PM_{2.5:10}$) or the count of particles measured in each size bin ($PM_{C0:1}$, $PM_{C1:2.5}$, $PM_{C2.5:10}$). Including T, gases other than CO, or alternative PM particulate size bins or counts had no noticeable effect on model performance, so these variables were ultimately omitted.

4.3.4 Evaluation of Generalizability and Drift. The Eclipse deployment framework uses results from the EPA co-location to remotely calibrate all other devices in the network. This approach assumes that calibration functions developed for co-located devices will generalize to other devices in other locations. To evaluate this assumption, we conducted “Leave-One-Device-Out” Cross-Validation (LODO CV), iteratively fitting models for all but one device and evaluating predictions for the excluded device. We also conducted “Leave-One-Station-Out” Cross-Validation (LOSO CV), iteratively excluding all devices at a given EPA station, for both $PM_{2.5}$ and O_3 . This approach enables an evaluation of the generalizability of our calibration function to other devices at locations other than those at which the function was parameterized, as is the case when the model is used to calibrate data across the entire network of devices.

We further note that the association between sensor readings and true values may drift over time. Although we did not consider drift to be a major concern given the relatively short observation

¹<https://gispub.epa.gov/airnow/>

²<https://www.epa.gov/aqs>

period reported here, annual drift may be significant: the manufacturers' stated estimates are $\pm 1.25 \mu\text{g}/\text{m}^3$ per year for $\text{PM}_{2.5}$, ± 100 ppb per year for CO, ± 20 ppb per year for O_3 , NO_2 , and SO_2 . We evaluated calibrated readings over periods when all three EPA stations report values close to zero, effectively providing a natural baseline for measurement. We also examined whether our calibration models introduce outliers by regressing the daily average calibrated and uncalibrated values from each sensor on those of all neighbors within 5km to evaluate whether any sensor produced notably different readings from its neighbors. We describe the results of these analyses in Section 6.3.

4.4 Participatory Network Design.

Local stakeholders hold expertise that can and should inform monitoring initiatives in their environments [16]. We shared our network design with local environmental justice organizations to identify areas missing coverage, and allocated a subset of devices ($N = 20$) to additional locations they selected. After the deployment, we facilitated recorded meetings, approved by our Institutional Review Board, to share data and to obtain feedback and analytics recommendations from each group. We also met regularly with city stakeholders to share network design, data access, and website details.

5 WEBSITE CREATION AND EVALUATION

To ensure that the data would be easily accessible to Chicago residents, city officials, and community partners, we created a website that visualizes air quality data in real time. (Data are also available via an API³). Two access points to the website each provide a different user view. First, the publicly accessible URL (<https://urban.microsoft.com/air/city/chicago>) brings the user to a website showing all of the Eclipse devices in Chicago with their most recent $\text{PM}_{2.5}$ readings (Fig. 6, left panel). Second, scanning a QR code posted at each bus station in the Eclipse network (Right panel of Fig. 6) brings users to an index page that explains the network and then to a website view that shows the prior week's data for that sensor. Both views allow users to select different locations and to examine a week of historical data.

We collected two distinct measures of the interactive experience. First, users who accessed the site via a QR code received the optional survey question "how does your air quality feel?". We included three possible responses (poor, fair, or good) to evaluate differences between perceived and measured air quality. Second, we collected anonymous telemetry data on website usage. We also conducted interviews with 11 city and community stakeholders to observe website usage and learn about additional or unmet analytic needs. This work was reviewed and approved by our organization's Institutional Review Board.

6 RESULTS

6.1 Spatio-temporal Variation and Reliability

Fig. 7 shows that the network detected citywide events such as wildfire plumes as well as hyperlocal pollution events. Over three months, we collected 221,737 sensor-hours of complete data. The

readings cover over 90% of expected sensor-hours of readings during this period, excluding data from the week of installations.

6.1.1 Device Maintenance. Of the 115 total devices deployed across the pilot, calibration, and main deployment stages, 44 devices required maintenance between the device's initial deployment date and September 30th. Key issues could be generally classified as follows and as shown in Fig. 8:

- Installation issues (5 devices). The device was not properly connected to the solar panel during the initial installation.
- Shelter issues (3 devices). Several problems stemmed from the structure of the bus shelter: one shelter was destroyed in a car accident, one shelter had routine maintenance scheduled, and the third shelter stopped getting enough sunlight in its downtown location after the end of the summer. These devices were all moved to new shelters.
- Signal issues (12 devices). We identified signal issues when devices disconnected from the LTE network, generally indicating insufficient signal strength for communication.
- Device issues (24 devices). This category encompassed the remaining hardware/software problems, such as overheated batteries or defective SIM cards.

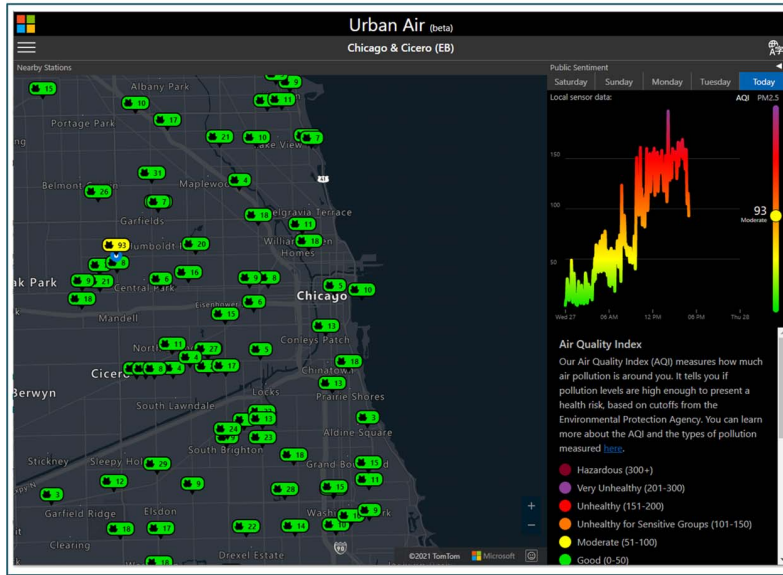
Installation and shelter issues were easily remedied by reconnecting solar panels or moving devices to the closest shelter. For signal and device issues, we used a "reset, relocate, replace" strategy. In this strategy, the first debugging step was to reset the device using a magnetic reset, avoiding the need for device removal. Of the 24 device issues, 14 were fixed with a reset, indicating a software issue. If resetting did not fix the issue, the device was relocated to a new bus shelter – a solution that resolved problems related to signal strength. Finally, if the issue persisted, the device was replaced.

Most issues (29 devices, 66%) surfaced within the first three weeks following installation. The vast majority of issues were also easily resolved by either resetting or relocating devices, highlighting the critical importance of monitoring and maintenance. The low per-device cost further enabled us to quickly replace faulty devices. As a measure of longevity, we note that through the full deployment as well as three months thereafter (July 2021 through December 2021), less than 10% of the total devices needed to be replaced.

6.1.2 Key Barriers for Urban IoT. Signal issues were common. Moving devices short distances (i.e. across the street) did not improve connectivity, but relocating to more distant locations fixed the issue in most cases. The distance to the closest cell tower was not associated with the likelihood of disconnecting, nor was there spatial clustering in the locations of the disconnected devices. These findings suggest that cellular dead zones, rather than a lack of coverage, may be impeding device connectivity. Urban form also impeded solar harvesting: at one location, tall buildings blocked incoming sunlight as the sun's angle decreased in September. Seasonal adapting of sampling and communication rates based on harvested-energy availability may be necessary to accommodate the unique limitations of urban settings. Finally, two devices were replaced due to destruction: one device was destroyed in a car accident; the other device was stolen from a community-selected site opposite a facility that the community had identified as a potential

³Details are available at <https://urban.microsoft.com/air/api/chicago/0?l=null>.

Urban Air WebApp for Public data access



QR-aided mobile data access



Figure 6: The air quality data is presented through a custom website (left). This website provides an interactive map view for the $PM_{2.5}$ data from every device as well as a line graph showing $PM_{2.5}$ levels over the previous week. Users can access the website by scanning the QR codes provided at augmented bus shelters with a mobile device (right).

source of air pollution. The community organization contacted the research team, and the stolen device was rapidly replaced.

6.2 Sensor Evaluation and Calibration

This section reports on the evaluation of the raw and calibrated Eclipse data in comparison with data from EPA monitoring stations. Calibration data were available from May 12 to September 30, 2021. During this period, regulatory monitors reported hourly $PM_{2.5}$ concentrations from 0 to $243.5 \mu\text{g}/\text{m}^3$, and hourly O_3 concentrations from 0 to 101 parts per billion (ppb). Hourly NO_2 concentrations over the reporting period ranged from 0.3 to 43.3 ppb and SO_2 concentrations ranged from 0 to 11.9 ppb. Eclipse devices reported hourly average relative humidity ranging from 12.2% to 95.7%.

6.2.1 Calibration Data Completeness. We cleaned the Eclipse data using a three-step QA/QC process. First, to mitigate the effects of outliers, gas data were winsorized (the top and bottom 0.1% of readings were trimmed to improve robustness to outliers) [28] and $PM_{2.5}$ data were dropped if they were outside manufacturer specifications ($PM_{2.5} = 0$ or $PM_{2.5} \geq 1000$, $N = 0$). Second, to exclude malfunctioning sensors, hourly data were excluded from sensor-hours with fewer than 9 readings ($N = 509$, 1.7%). Third, daily data were excluded from sensor-days with fewer than 18 readings ($N = 186$, 0.6%). One device malfunctioned and reported only intermittently; it was excluded from the analysis and replaced with a new device on July 15, resulting in the loss of 1,560 expected sensor-hours of data (5.1%). After QA/QC, the co-located devices captured 92.6% of expected hourly observations ($N = 28,417$).

6.2.2 Calibration Model Selection. The final calibration models were selected based on performance (evaluated with R^2). Hourly

$PM_{2.5}$ was modeled as a function of hourly Eclipse $PM_{2.5}$ and CO readings, with relative humidity adjusted for by including both B-splines with four knots and allowing an interaction of a dummy variable equal to 1 when $RH > 88\%$ with Eclipse $PM_{2.5}$ – an adjustment that is similar to the approaches used by Zusman et al. [64]. For gases, random forest models performed significantly better than other models, consistent with prior research [15, 63]; final models adjusted for temperature, relative humidity, zero offsets, and the full array of other pollutant readings from the Eclipse device.

6.2.3 Accuracy and Precision. Fig. 9 shows the precision and accuracy of each sensor evaluated with raw data, cross-validated with training data, or evaluated on the test set. Notably, the raw $PM_{2.5}$ data are precise ($CV = 11.8\%$) but inaccurate ($NRMSE = 49.3\%$). This level of error is similar to levels attained by other low-cost sensors, such as PurpleAir sensors [4], and EPA notes that even these relatively inaccurate data can be useful for education and information [62]. However, calibration reduces $NRMSE$ by approximately 41.2% and reduces CV by 44.1%. The calibrated data thus achieve levels of accuracy recommended for personal exposure or hotspot identification and characterization ($NRMSE < 0.3$) [62]. Moreover, the test set $RMSE$ (3.1) is comparable to the results of EPA calibration models for PurpleAir sensors ($RMSE = 3$) [4].

The gas sensors perform more poorly in the field. Raw O_3 is imprecise ($CV = 219\%$) and inaccurate ($NRMSE = 113\%$). Raw NO_2 and SO_2 data are moderately precise ($CV = 18.5\%$ and $CV = 33.5\%$, respectively) but include errors larger than the average EPA readings by a factor of 6 for NO_2 and a factor of 43 for SO_2 . Calibration contributes large reductions in error – $NRMSE$ is reduced by 73.6% for O_3 , 88.6% for NO_2 , and 98.8% for SO_2 – improving precision to

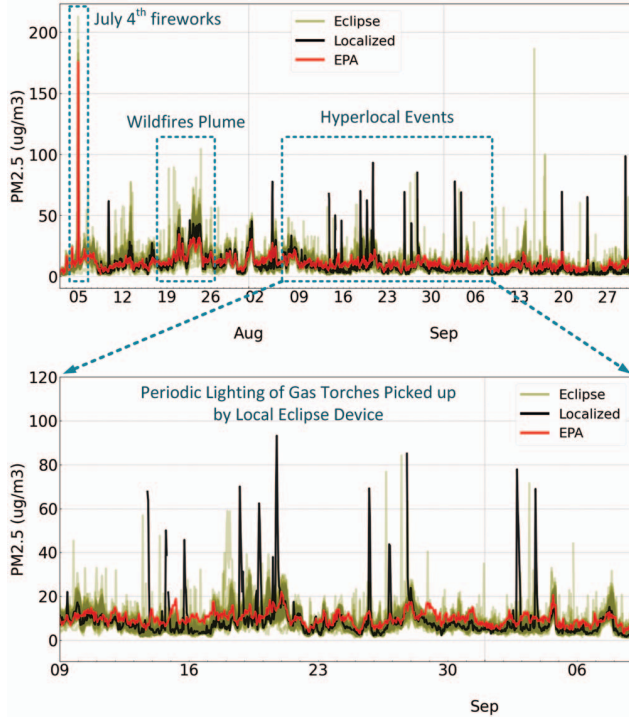


Figure 7: Fine Particulate Matter ($PM_{2.5}$) during the three month deployment. Green lines show raw hourly average $PM_{2.5}$ at each Eclipse Device; the red line indicates average hourly values at EPA regulatory monitors. The top panel highlights two major pollution events that were detected by both EPA and the Eclipse sensors (July 4 and wildfires) as well as hyperlocal events (gas torches [29]). The bottom panel shows periodic spikes from 08/09 and 09/09 that were only observed in Eclipse nodes closest to the source.




Figure 8: Breakdown of the 44 different device issue types that arose from device deployment until September 30, 2021.

levels consistent with guidelines for supplementary monitoring; error remains relatively high, however, and the final NRMSE for calibrated NO_2 (0.71) exceeds that of recent results from the Breathe London campaign (0.35) [49].

Both the calibrated $PM_{2.5}$ and the calibrated O_3 data exhibit linear relationships, in the test set, with gold standard EPA data




Figure 9: Precision (top) and error (bottom) of Eclipse devices in comparison with EPA data. Dashed lines indicate the thresholds recommended by EPA for A) Education and Information (0.5); B) Hotspot Identification and Characterization or Personal Exposure Monitoring (0.3); C) Supplemental Monitoring (0.2); and D) Regulatory Monitoring (0.1).

(slope = 0.75 and intercept = 2.57 for $PM_{2.5}$; slope = 0.97 and intercept = 0.23 for O_3). Calibrated $PM_{2.5}$ approaches the EPA's recommended threshold for R^2 (0.7) [22, 62] although calibrated O_3 remains slightly lower ($R^2 = 0.66$ in comparison with EPA recommendation of 0.8). Measures are poorer for NO_2 ($R^2 = 0.29$) and SO_2 ($R^2 = 0.42$), a result likely attributable to low signal-to-noise ratios due to the generally low levels of ambient NO_2 and SO_2 at EPA sites [43]. These results represent small but notable improvements over early results from Array of Things (O_3 $R^2=0.59$ and NO_2 $R^2 = 0.14$) [50].

6.2.4 *Model Generalizability and Evidence of Drift.* Three findings suggest that the calibration models should generalize from the co-located sensors to the network as a whole. First, the moderate precision of raw data and high precision of calibrated data suggest that calibration conducted with one sensor should generalize to another device. Second, the consistency of LODO CV with 5-Fold CV for NRMSE on the test set further supports the assumption of generalizability of calibration models to an excluded device. Finally, LOSO CV shows that calibration models for $PM_{2.5}$ and O_3 perform as expected at excluded stations.




Figure 10: There were over 800 unique website sessions over the course of the three-month deployment, but less than 15% of sessions accessed the website by scanning a QR code.

We do see evidence of drift for $PM_{2.5}$, O_3 , and NO_2 when comparing night-time readings from our sensors and regulatory monitors in September versus July (mean $1.6 \mu g/m^3$, 13.2 ppb, and -12.4 ppb, respectively); after calibration, these differences are attenuated but still significant ($PM_{2.5}$: 0.7 95% CI 0.4 - 1.0; O_3 : 2.3 95% CI 0.4 - 4.1; NO_2 : 1.1 95% CI 1.0-1.2). Sensors generally produce similar daily average readings to neighbors within 5km when using uncalibrated and calibrated readings: we observe $R^2 > 0.5$ for all but two sensors for calibrated $PM_{2.5}$ and for all sensors using calibrated O_3 . The explanatory power of neighboring sensors was lower for calibrated versus uncalibrated NO_2 (R^2 IQR 0.35 - 0.57 versus 0.68 - 0.93) and SO_2 (R^2 IQR 0.42 - 0.60 versus 0.47 - 0.84) but, given the tighter range of the calibrated data, these changes likely do not reflect the introduction of outliers but rather support the ability of calibrated models to capture true hyperlocal variation, which is known to be larger for NO_2 versus for $PM_{2.5}$ or O_3 [12].

6.3 Website Usage

The website offered a convenient and accessible platform for sharing results. Over the first three months of the full deployment (July through September, 2021), there were over 800 unique sessions. However, Fig. 10 shows that access via QR codes was limited. Individual QR codes were scanned a total of 113 times between July 1 and September 30th, representing less than 15% of unique sessions. (To protect anonymity, this project did not use cookies; we thus used the number of unique device-day combinations as a lower bound for the number of sessions.) This result suggests that a QR code is a barrier impeding data access; a solution may be to further augment the sensors or bus shelters with visual or auditory experiences – that is, removing the barrier between users and data entirely. Local partners noted that awareness of the data (and the sticker with the QR code) might also be limiting access, illustrating the need for ongoing partnerships and communication with community groups. Nevertheless, for those people who did scan the QR code and answer the user survey ($N = 68$), perceived and actual air quality were correlated: we recorded significantly higher $PM_{2.5}$ levels for locations and times where users reported “poor” or “fair” air versus “good” air ($p < 0.01$).

7 DISCUSSION

This paper describes the development and implementation of a large-scale, multi-pollutant IoT monitoring network across a major

North American city. We provide an end-to-end deployment framework, including: (1) the design of a low-cost device that improves upon previous research via its compact, aesthetic structure and its modular sensing capability, making it easy to tailor to and deploy in different cities; (2) a deployment strategy that piggybacks on existing infrastructure to minimize installation and maintenance effort and costs; (3) a concurrent co-location with EPA regulatory monitors to improve accuracy via ongoing calibration; and (4) cloud analytics including a website that provides data to the public in real time. Although this application focused on air quality sensing, our approach provides a replicable and scalable template that could be adapted to other urban sensing networks to measure environmental exposures such as noise or heat sensing.

To support similar initiatives, we reflect on key lessons learned from our experience with a large-scale, real-world deployment and share the ongoing needs gleaned from our network of community and city stakeholders.

- (1) *Urban environments pose unique challenges.* Tall buildings can obstruct communications signals. To overcome cellular dead zones, networks may need to incorporate repeaters or other solutions to amplify device signals. Building shadows also limit the locations where solar-powered devices can harvest adequate energy without requiring trade-offs in sampling frequency.
- (2) *Low-cost sensors are a complement, not a substitute, to regulatory networks.* Ongoing co-locations are needed for calibration, and network designers should consider budgeting for a regulatory-grade sensor when constructing similar networks in Global South cities or other places that lack established regulatory networks.
- (3) *Maintenance and monitoring are key.* Minimizing hardware costs, so that faulty or stolen devices could easily be replaced, and ensuring that resetting and relocating were part of city or partner employees’ existing daily procedures, was critical to minimize data loss.
- (4) *Involve local stakeholders early and often.* City and community leaders have expertise such as knowledge of potential polluters and insights on policy priorities; these stakeholders’ insights informed our selection of gas sensors and enabled us to identify gaps in the network in advance of deployment in ways that made the resulting network more relevant and usable for local priorities such as the development of a cumulative hazards ordinance.

These lessons highlight opportunities for future work in low-cost urban sensing networks. Research is needed that evaluates the use of repeaters or to further develop new types of networks (e.g. LoRaWan) for improving connectivity in built-up areas, but also finds methods to minimize the additional installation and maintenance costs. Further work should identify power-aware sampling approaches that minimize data loss and maintain accuracy in places where buildings or other elements of the urban form limit sunlight. Further study is also needed to improve calibration models, both to boost signal-to-noise ratios for electrochemical gas sensors and to develop field calibration solutions that do not rely on regulatory co-locations for implementation in resource-constrained settings.

Finally, our work relied on random sampling and participatory network design to ensure that our deployment could provide insights both on differences between neighborhoods and on particular locations of local concern; however, there remains a need for studies to evaluate the optimal number and density of sensors needed for a given network. Because the solution to such a problem likely varies based on city size and morphology, pollutant examined, and the purpose of the monitoring network, such research requires ongoing exploratory network deployments across global cities.

8 ACKNOWLEDGEMENTS

We would like to thank Raed Mansour, Meera Raja, Jamie Ponce, and Adam Heckman, as well as from Tiffany Werner and the Environmental Law and Policy Center, for their guidance. This work was also made possible by support from Pinaki Banerjee and Rajesh Sankaran, as well as Gabrielle Brussel, Nicolas Clochard-Bossuet, Pavol Rehus, and the JCDecaux leadership and maintenance teams. Tony Carbury, Precious Esie, and Nneka Udeagbala contributed research support. Finally, we are grateful for insights from BUILD, Inc., Little Village Environmental Justice Organization, Neighbors 4 Environmental Justice, Blacks in Green, and Latinos Progresando.

REFERENCES

- [1] ADKINS, J., GHENA, B., JACKSON, N., PANNUTO, P., ROHRER, S., CAMPBELL, B., AND DUTTA, P. The signpost platform for city-scale sensing. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)* (2018), IEEE, pp. 188–199.
- [2] ANJOMSHOAA, A., DUARTE, F., RENNINGS, D., MATARAZZO, T. J., DESOUSA, P., AND RATTI, C. City scanner: Building and scheduling a mobile sensing platform for smart city services. *IEEE Internet of Things Journal* 5, 6 (2018), 4567–4579.
- [3] APTE, J. S., MESSIER, K. P., GANI, S., BRAUER, M., KIRCHSTETTER, T. W., LUNDEN, M. M., MARSHALL, J. D., PORTIER, C. J., VERMEULEN, R. C., AND HAMBURG, S. P. High-resolution air pollution mapping with Google Street View cars: Exploiting big data. *Environmental Science & Technology* 51, 12 (2017), 6999–7008.
- [4] BARKJOHN, K. K., GANTT, B., AND CLEMENTS, A. L. Development and application of a United States-wide correction for PM 2.5 data collected with the PurpleAir sensor. *Atmospheric Measurement Techniques* 14, 6 (2021), 4617–4637.
- [5] BECNEL, T., TINGEY, K., WHITAKER, J., SAYAHI, T., LÉ, K., GOFFIN, P., BUTTERFIELD, A., KELLY, K., AND GAILLARDON, P.-E. A distributed low-cost pollution monitoring platform. *IEEE Internet of Things Journal* 6, 6 (2019), 10738–10748.
- [6] BI, J., WILDANI, A., CHANG, H. H., AND LIU, Y. Incorporating low-cost sensor measurements into high-resolution PM2.5 modeling at a large spatial scale. *Environmental Science & Technology* 54, 4 (2020), 2152–2162.
- [7] Breathe london technical report pilot phase (2018 – 2020). <https://www.globalcleanair.org/files/2021/05/BL-Pilot-Final-Technical-Report.pdf>, 2021.
- [8] CARVLIN, G. N., LUGO, H., OLMEDO, L., BEJARANO, E., WILKIE, A., MELTZER, D., WONG, M., KING, G., NORTHCROSS, A., JERRETT, M., ET AL. Development and field validation of a community-engaged particulate matter air quality monitoring network in Imperial, California, USA. *Journal of the Air & Waste Management Association* 67, 12 (2017), 1342–1352.
- [9] CASTELL, N., DAUGE, F. R., SCHNEIDER, P., VOGT, M., LERNER, U., FISHBAIN, B., BRODAY, D., AND BARTONOVA, A. Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates? *Environment International* 99 (2017), 293–302.
- [10] CATLETT, C., BECKMAN, B., FERRIER, N., NUSBAUM, H., PAPKA, M., BERMAN, M., AND SANKARAN, R. Measuring cities with software-defined sensors. *Journal of Social Computing* 1, 1 (2020).
- [11] CATLETT, C., BECKMAN, P., SANKARAN, R., AND GALVIN, K. K. Array of Things: A scientific research instrument in the public way: Platform design and early lessons learned. In *Proceedings of the 2nd International Workshop on Science of Smart City Operations and Platforms Engineering* (2017), ACM, pp. 26–33.
- [12] CHAMBLISS, S., PINON, C., MESSIER, K., LAFRANCHI, B., UPPERMAN, C., LUNDEN, M., ROBINSON, A., MARSHALL, J. D., AND APTE, J. Local and regional-scale racial and ethnic disparities in air pollution determined by long-term mobile monitoring. *Proceedings of the National Academy of Sciences* (2021).
- [13] CHEN, L.-J., HO, Y.-H., LEE, H.-C., WU, H.-C., LIU, H.-M., HSIEH, H.-H., HUANG, Y.-T., AND LUNG, S.-C. C. An open framework for participatory PM2.5 monitoring in smart cities. *IEEE Access* 5 (2017), 14441–14454.
- [14] CHEN, M., YANG, J., HU, L., HOSSAIN, M. S., AND MUHAMMAD, G. Urban healthcare big data system based on crowdsourced and cloud-based air quality indicators. *IEEE Communications Magazine* 56, 11 (2018), 14–20.
- [15] CONCAS, F., MINERAUD, J., LAGERSPETZ, E., VARJONEN, S., LIU, X., PUOLAMÄKI, K., NURMI, P., AND TARKOMA, S. Low-cost outdoor air quality monitoring and sensor calibration: A survey and critical analysis. *ACM Transactions on Sensor Networks (TOSN)* 17, 2 (2021), 1–44.
- [16] CORBURN, J. *Street science: Community knowledge and environmental health justice*. MIT Press, 2005.
- [17] CROSS, E. S., WILLIAMS, L. R., LEWIS, D. K., MAGOON, G. R., ONASCH, T. B., KAMINSKY, M. L., WORSNOP, D. R., AND JAYNE, J. T. Use of electrochemical sensors for measurement of air pollution: correcting interference response and validating measurements. *Atmospheric Measurement Techniques* 10, 9 (2017), 3575–3588.
- [18] DAEPP, M. I. G., BINET, A., GAVIN, V., ARCAJA, M. C., AND THE HEALTHY NEIGHBORHOODS RESEARCH CONSORTIUM. The moving mapper: Participatory action research with big data. *Journal of the American Planning Association* (2021), 1–13.
- [19] DAVIDSON, C. I., PHALEN, R. F., AND SOLOMON, P. A. Airborne particulate matter and human health: a review. *Aerosol Science and Technology* 39, 8 (2005), 737–749.
- [20] DESOUSA, P., ANJOMSHOAA, A., DUARTE, F., KAHN, R., KUMAR, P., AND RATTI, C. Air quality monitoring using mobile low-cost sensors mounted on trash-trucks: Methods development and lessons learned. *Sustainable Cities and Society* 60 (2020), 102239.
- [21] DESOUSA, P., AND KINNEY, P. L. On the distribution of low-cost PM 2.5 sensors in the US: demographic and air quality associations. *Journal of Exposure Science & Environmental Epidemiology* 31, 3 (2021), 514–524.
- [22] DUVAL, R., CLEMENTS, A., HAGLER, G., KAMAL, A., KILARU, V., GOODMAN, L., FREDERICK, S., BARKJOHN, K. J., VONWALD, I., GREENE, D., AND DYE, T. Performance testing protocols, metrics, and target values for fine particulate matter air sensors: Use in ambient, outdoor, fixed site, non-regulatory supplemental and informational monitoring applications. https://cfpub.epa.gov/si/si_public_record_report.cfm?Lab=CEMM&dirEntryId=350785, 2021.
- [23] DUVAL, R., CLEMENTS, A., HAGLER, G., KAMAL, A., KILARU, V., GOODMAN, L., FREDERICK, S., BARKJOHN, K. J., VONWALD, I., GREENE, D., AND DYE, T. Performance testing protocols, metrics, and target values for ozone air sensors: Use in ambient, outdoor, fixed site, non-regulatory and informational monitoring applications. https://cfpub.epa.gov/si/si_public_record_report.cfm?Lab=CEMM&dirEntryId=350784, 2021.
- [24] GREEN, B. *The smart enough city: putting technology in its place to reclaim our urban future*. MIT Press, 2019.
- [25] HAGAN, D. H., ISAACMAN-VANWERTZ, G., FRANKLIN, J. P., WALLACE, L. M., KOCAR, B. D., HEALD, C. L., AND KROLL, J. H. Calibration and assessment of electrochemical air quality sensors by co-location with regulatory-grade instruments. *Atmospheric Measurement Techniques* 11, 1 (2018), 315–328.
- [26] HAN, P., MEI, H., LIU, D., ZENG, N., TANG, X., WANG, Y., AND PAN, Y. Calibrations of low-cost air pollution monitoring sensors for CO, NO2, O3, and SO2. *Sensors* 21, 1 (2021), 256.
- [27] HANKEY, S., AND MARSHALL, J. D. On-bicycle exposure to particulate air pollution: Particle number, black carbon, PM 2.5, and particle size. *Atmospheric Environment* 122 (2015), 65–73.
- [28] HASTINGS JR, C., MOSTELLER, F., TUKEY, J. W., AND WINSOR, C. P. Low moments for small samples: a comparative study of order statistics. *The Annals of Mathematical Statistics* 18, 3 (1947), 413–426.
- [29] HECKTMAN, A. J. Every datum tells a story, don't it? <https://chitechforgood.com/f/every-datum-tells-a-story-dont-it>, 2021.
- [30] HERTWIG, D., GRIMMOND, S., KOTTHAUS, S., VANDERWEL, C., GOUGH, H., HAEFELIN, M., AND ROBINS, A. Variability of physical meteorology in urban areas at different scales: implications for air quality. *Faraday Discussions* 226 (2021), 149–172.
- [31] HOEK, G., BEELEN, R., DE HOOGH, K., VIENNEAU, D., GULLIVER, J., FISCHER, P., AND BRIGGS, D. A review of land-use regression models to assess spatial variation of outdoor air pollution. *Atmospheric environment* 42, 33 (2008), 7561–7578.
- [32] HUANG, J., DUAN, N., JI, P., MA, C., DING, Y., YU, Y., ZHOU, Q., SUN, W., ET AL. A crowdsourcing-based sensing system for monitoring fine-grained air quality in urban environments. *IEEE Internet of Things Journal* 6, 2 (2018), 3240–3247.
- [33] ILLINOIS ENVIRONMENTAL PROTECTION AGENCY. Illinois ambient air monitoring 2021 network plan. <https://www2.illinois.gov/epa/topics/air-quality/outdoor-air/air-monitoring/Documents/2021%20Network%20Plan.pdf>, 2020.
- [34] KELLY, K. E., XING, W. W., SAYAHI, T., MITCHELL, L., BECNEL, T., GAILLARDON, P.-E., MEYER, M., AND WHITAKER, R. T. Community-based measurements reveal unseen differences during air pollution episodes. *Environmental Science & Technology* 55, 1 (2020), 120–128.
- [35] KHEIRBEK, I., HANEY, J., DOUGLAS, S., ITO, K., CAPUTO JR, S., AND MATTE, T. The public health benefits of reducing fine particulate matter through conversion to cleaner heating fuels in New York City. *Environmental Science & Technology* 48, 23 (2014), 13573–13582.
- [36] KIM, K.-H., KABIR, E., AND KABIR, S. A review on the human health impact of airborne particulate matter. *Environment International* 74 (2015), 136–143.

- [37] KRIEGER, N., WATERMAN, P. D., BATRA, N., MURPHY, J. S., DOOLEY, D. P., AND SHAH, S. N. Measures of local segregation for monitoring health inequities by local health departments. *American Journal of Public Health* 107, 6 (2017), 903–906.
- [38] LEE, H., KANG, J., KIM, S., IM, Y., YOO, S., AND LEE, D. Long-term evaluation and calibration of low-cost particulate matter (pm) sensor. *Sensors* 20, 13 (2020), 3617.
- [39] LEWIS, A., AND EDWARDS, P. Validate personal air-pollution sensors. *Nature* 535, 7610 (2016), 29.
- [40] LEWIS, A., PELTIER, W. R., AND VON SCHNEIDEMESSER, E. Low-cost sensors for the measurement of atmospheric composition: overview of topic and future applications. World Meteorological Organization (WMO), 2018.
- [41] LI, J. J., FALTINGS, B., SAUKH, O., HASENFRATZ, D., AND BEUTEL, J. Sensing the air we breathe—the opensense Zurich dataset. In *Twenty-Sixth AAAI Conference on Artificial Intelligence* (2012).
- [42] LIANG, Y., SENGUPTA, D., CAMPPIER, M. J., APTE, J., LUNDBERG, D. M., AND GOLDSTEIN, A. Wildfire smoke impacts on indoor air quality assessed using crowdsourced data in California. *Proceedings of the National Academy of Sciences* (2021).
- [43] MAAG, B., ZHOU, Z., AND THIELE, L. A survey on sensor calibration in air pollution monitoring deployments. *IEEE Internet of Things Journal* 5, 6 (2018), 4857–4870.
- [44] MATTE, T. D., ROSS, Z., KHEIRBEK, I., EISEL, H., JOHNSON, S., GORCZYNSKI, J. E., KASS, D., MARKOWITZ, S., PEZESHKI, G., AND CLOUGHERTY, J. E. Monitoring intraurban spatial patterns of multiple combustion air pollutants in New York City: design and implementation. *Journal of exposure science & environmental epidemiology* 23, 3 (2013), 223–231.
- [45] MOLTCHANOV, S., LEVY, I., ETZION, Y., LERNER, U., BRODAY, D. M., AND FISHBAIN, B. On the feasibility of measuring urban air pollution by wireless distributed sensor networks. *Science of The Total Environment* 502 (2015), 537–547.
- [46] MORAWSKA, L., THAI, P. K., LIU, X., ASUMADU-SAKYI, A., AYOKO, G., BARTONOVA, A., BEDINI, A., CHAI, F., CHRISTENSEN, B., DUNBABIN, M., ET AL. Applications of low-cost sensing technologies for air quality monitoring and exposure assessment: How far have they gone? *Environment International* 116 (2018), 286–299.
- [47] MOUSAVI, A., AND WU, J. Indoor-generated PM_{2.5} during COVID-19 shutdowns across California: application of the PurpleAir indoor–outdoor low-cost sensor network. *Environmental Science & Technology* 55, 9 (2021), 5648–5656.
- [48] PAAS, B., SCHMIDT, T., MARKOVA, S., MARAS, I., ZIEFLE, M., AND SCHNEIDER, C. Small-scale variability of particulate matter and perception of air quality in an inner-city recreational area in Aachen, Germany. *Meteorologische Zeitschrift* (2016), 305–317.
- [49] PETERS, D. R., POPOOLA, O. A. M., JONES, R. L., MARTIN, N. A., MILLS, J., FONSECA, E. R., STIDWORTHY, A., FORSYTH, E., CARRUTHERS, D., DUPUY-TODD, M., DOUGLAS, F., MOORE, K., SHAH, R. U., PADILLA, L. E., AND ALVAREZ, R. A. Evaluating uncertainty in sensor networks for urban air pollution insights. *Atmospheric Measurement Techniques* 15, 2 (2022), 321–334.
- [50] POTOSNAK, M., BANERJEE, P., SANKARAN, R., KOTAMARTHI, V., JACOB, R., BECKMAN, P., AND CATLETT, C. Array of things: Characterizing low-cost air quality sensors for a city-wide instrument. *Earth and Space Science Open Archive* (2018), 1.
- [51] SAYAHI, T., BUTTERFIELD, A., AND KELLY, K. Long-term field evaluation of the Plantower PMS low-cost particulate matter sensors. *Environmental pollution* 245 (2019), 932–940.
- [52] SHUSTERMAN, A. A., TEIGE, V. E., TURNER, A. J., NEWMAN, C., KIM, J., AND COHEN, R. C. The Berkeley Atmospheric CO₂ Observation Network: Initial evaluation. *Atmospheric Chemistry and Physics* 16, 21 (2016), 13449–13463.
- [53] SOUTH COAST AIR QUALITY MANAGEMENT DISTRICT. Field evaluation: Sensirion SPS30. <http://www.aqmd.gov/docs/default-source/air-spec/field-evaluations/sensirion-sps30-evaluation-kit--field-evaluation.pdf?sfvrsn=42>, 2019.
- [54] TANZER, R., MALINGS, C., HAURYLIUK, A., SUBRAMANIAN, R., AND PRESTO, A. A. Demonstration of a low-cost multi-pollutant network to quantify intra-urban spatial variations in air pollutant source impacts and to evaluate environmental justice. *International Journal of Environmental Research and Public Health* 16, 14 (2019), 2523.
- [55] THORNTON, S. A guide to Chicago’s Array of Things Initiative. *Ash Center for Democratic Governance and Innovation, Harvard Kennedy School*. (2018). <https://datasmart.ash.harvard.edu/news/article/a-guide-to-chicagos-array-of-things-initiative-1190>.
- [56] TRYNER, J., MEHAFFY, J., MILLER-LIONBERG, D., AND VOLCKENS, J. Effects of aerosol type and simulated aging on performance of low-cost PM sensors. *Journal of Aerosol Science* 150 (2020), 105654.
- [57] TURNER, A. J., SHUSTERMAN, A. A., McDONALD, B. C., TEIGE, V., HARLEY, R. A., AND COHEN, R. C. Network design for quantifying urban CO₂ emissions: assessing trade-offs between precision and network density. *Atmospheric Chemistry and Physics* 16, 21 (2016), 13465–13475.
- [58] UNITED NATIONS. 68% of the world population projected to live in urban areas by 2050, says UN. <https://www.un.org/sustainabledevelopment/blog/2018/05/68-of-the-world-population-projected-to-live-in-urban-areas-by-2050-says-un/>, 2018.
- [59] UNITED STATES ENVIRONMENTAL PROTECTION AGENCY. National Ambient Air Quality Standards. <https://www.epa.gov/criteria-air-pollutants/naaqs-table>.
- [60] VOS, T., LIM, S. S., ABBAFATI, C., ABBAS, K. M., ABBASI, M., ABBASIFARD, M., ABBASI-KANGEVARI, M., ABBASTABAR, H., ABD-ALLAH, F., ABDELALIM, A., ET AL. Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the global burden of disease study 2019. *The Lancet* 396, 10258 (2020), 1204–1222.
- [61] WANG, M., AARON, C. P., MADRIGANO, J., HOFFMAN, E. A., ANGELINI, E., YANG, J., LAINE, A., VETTERLI, T. M., KINNEY, P. L., SAMPSON, P. D., ET AL. Association between long-term exposure to ambient air pollution and change in quantitatively assessed emphysema and lung function. *JAMA* 322, 6 (2019), 546–556.
- [62] WILLIAMS, R., KILARU, V., SNYDER, E., KAUFMAN, A., DYE, T., RUTTER, A., RUSSELL, A., AND HAFNER, H. Air sensor guidebook. https://cfpub.epa.gov/si/si_public_record_report.cfm?Lab=NERL&dirEntryId=277996&simpleSearch=1&searchAll=air+sensor+guidebook, 2014.
- [63] ZIMMERMAN, N., PRESTO, A. A., KUMAR, S. P., GU, J., HAURYLIUK, A., ROBINSON, E. S., ROBINSON, A. L., AND SUBRAMANIAN, R. A machine learning calibration model using random forests to improve sensor performance for lower-cost air quality monitoring. *Atmospheric Measurement Techniques* 11, 1 (2018), 291–313.
- [64] ZUSMAN, M., SCHUMACHER, C. S., GASSETT, A. J., SPALT, E. W., AUSTIN, E., LARSON, T. V., CARVLIN, G., SETO, E., KAUFMAN, J. D., AND SHEPPARD, L. Calibration of low-cost particulate matter sensors: Model development for a multi-city epidemiological study. *Environment International* 134 (2020), 105329.