PortoLivingLab: An IoT-Based Sensing Platform for Smart Cities

Pedro M. Santos, João G. P. Rodrigues, Susana B. Cruz, Tiago Lourenço, Pedro M. d'Orey, Yunior Luis, Cecília Rocha, Sofia Sousa, Sérgio Crisóstomo, Cristina Queirós, Susana Sargento, *Member, IEEE*, Ana Aguiar, *Member, IEEE*, and João Barros, *Senior Member, IEEE*

Abstract—**Smart cities aim to improve the citizens' quality of life by leveraging information about urban scale processes extracted from heterogeneous data sources collected on city-wide deployments. The Internet-of-Things (IoT) is, thus, the enabler of smart city technologies at urban scale. In this paper, we present PortoLivingLab, a multisource sensing infrastructure that leverages IoT technology to achieve city-scale sensing of four phenomena: weather, environment, public transport, and people flows. To sense these processes on a city scale, we deployed a vehicular network with over 600 vehicles and 19 static environmental sensors. We also developed an easily reconfigurable crowdsensing platform and carried out several crowdsensing campaigns with more than 600 participants. The data is collected in a common backend and stored using similar spatio-temporal data models to simplify sharing and joint analysis for the characterization of urban dynamics. We describe the architecture and composing elements of PortoLivingLab, highlighting the IoT technologies, and challenges faced. We present several proof-of-concept use cases (e.g., passenger flows from WiFi connections) that provide new insights into different components of an evolving and moving city. Finally, we lay out the future lines of work that will strive for finding hidden phenomena by leveraging data from the three complementary platforms.**

Index Terms—**Internet-of-Things (IoT), smart cities, urban sensing.**

I. INTRODUCTION

CITIES are complex systems that encompass a large number of simultaneous underlying processes at various levels, ranging from social systems (e.g., people flows), services (e.g., as transport), resources (e.g., air quality), and natural environment (e.g., as weather). They are known to have short-term variations, but also to show steady underlying patterns. Smart cities [1] aim at improving the citizens' quality of life by providing a deeper understanding of the urban processes, which in turn supports better-informed policy decisions and make new services available to the citizens based on that acquired knowledge. The variety of the processes to observe requires sourcing heterogeneous technologies for sensing and data collection, whereas deployment requirements in terms of scale, geographic spread and incurred costs demand for efficient technology usage strategies. The Internet-of-Things (IoT) [2] paradigm encompasses the tools and solutions for such large-scale, pervasive, and distributed monitoring systems. The ultimate goal is to build descriptive and predictive models for the phenomena of interest. As these phenomena are interdependent, it will be necessary to combine and fuse data from the different sensor sources.

In this paper, we share our learn-by-doing experience[1] in building a city-scale IoT deployment to monitor the weather, environment (air quality and noise), and mobility. The first contribution of this paper is the definition of the architecture of a city-scale sensing system deployed to observe relevant variables for each phenomena with static, mobile, and personal sensors. Then, we show how we combined a variety of backhaul technologies to retrieve the sensed data to a backend. We also provide additional information on the server backend components focusing on data management, processing and sharing issues. Finally, we discuss challenges and showcase use cases of how collected data may be used to learn or quantify more precisely urban phenomena.

PortoLivingLab is an urban-scale, multisource sensing infrastructure deployed in the city of Porto, Portugal. This infrastructure builds on three monitoring platforms.

1) *SenseMyCity (SMC) [3]:* A mobile crowdsensing research tool for gathering data from a large number of

Manuscript received December 31, 2016; revised August 30, 2017 and December 7, 2017; accepted December 24, 2017. Date of publication January 9, 2018; date of current version April 10, 2018. This work was supported in part by the European FP7 Project Future Cities under Grant FP7-REGPOT-2012-2013-1, in part by the Projects SenseBusNet and SmartCityMules at the Instituto de Telecomunicações, in part by the Research and Development Unit 50008 through applicable financial framework (FCT/MEC through national funds and FEDER-PT2020 partnership agreement when applicable) under Grant UID/EEA/50008/2013, and in part by the CMU-Portugal Program through "S2MovingCity: Sensing and Serving a Moving City" under Grant CMUPERI/TIC/0010/2014. *(Corresponding author: Pedro M. Santos.)*

P. M. Santos, T. Lourenço, Y. Luis, and J. Barros are with the Faculty of Engineering, Department of Electrical and Computer Engineering, University of Porto, 4099-002 Porto, Portugal (e-mail: pmsantos@fe.up.pt; up201402609@fe.up.pt; jrluis@fe.up.pt; jbarros@fe.up.pt).

J. G. P. Rodrigues, S. B. Cruz, P. M. d'Orey, S. Crisóstomo, and A. Aguiar are with the University of Porto, 4099-002 Porto, Portugal, and also with the Instituto de Telecomunicações, 4200-465 Porto, Portugal (e-mail: joao.g.p.rodrigues@fe.up.pt; susana.bulas.cruz@fe.up.pt; pedro.dorey@dcc.fc.up.pt; slcrisos@fc.up.pt; anaa@fe.up.pt).

C. Rocha is with the Faculty of Engineering, Department of Civil Engineering, University of Porto, 4099-002 Porto, Portugal (e-mail: carocha@fe.up.pt).

S. Sousa is with the Laboratório de Engenharia de Processos, Ambiente, Biotecnologia e Energia, Faculty of Engineering, University of Porto, 4099-002 Porto, Portugal (e-mail: sisousa@fe.up.pt).

C. Queirós is with the Faculty of Psychology and Sciences of Education, University of Porto, 4099-002 Porto, Portugal (e-mail: cqueiros@fpce.up.pt).

S. Sargento is with the Department of Electronics, Telecommunications and Informatics Engineering, University of Aveiro, 3810-193 Aveiro, Portugal, and also with the Instituto de Telecomunicações, 3810-193 Porto, Portugal (e-mail: susana@ua.pt).

Digital Object Identifier 10.1109/JIOT.2018.2791522

---

[1] [Online]. Available: https://pages.nist.gov/smartcitiesarchitecture

participants and their surroundings using a smartphone application.

2) *UrbanSense [4]:* A set of low-cost sensing stations deployed in 19 strategic locations to monitor environmental parameters (as noise and air pollution) and weather conditions.
3) *BusNet [5]:* A sensing and data distribution platform composed by 600+ public transport vehicles equipped with on-board units (OBUs) for V2X communications and in-vehicle Internet connectivity for passengers.

The three platforms have been operating in parallel for almost two years. All platforms have high levels of connectivity that enable IoT-inspired data collection strategies. The datasets acquired by PortoLivingLab have been used to describe the urban dynamics, by means of spatio-temporal maps of speed, origin–destination matrices with mode information, WiFi hotspot density, traffic state estimation, among other examples.

The remainder of this paper is organized as follows. Section II reviews existing literature on IoT architectures and urban-scale platforms. Section III presents the architecture of the city-scale data gathering platform. In Section IV, we detail the scope and operation of the three urban sensing platforms deployed in the city of Porto. Section V describes the communication backhauls and protocols of the platform, while the backend infrastructure is detailed in Section VI. Section VII summarizes the platform characteristics, and showcases use case studies. Final remarks are given in Section VIII.

II. RELATED WORK

Jin *et al.* [6] proposes a dataflow-driven description of the architecture and operation of IoT platforms for smart cities, composed of the stages *collection*, *processing*, *management*, and *interpretation*. The survey [7] proposes a list of design and operation planes of an IoT platforms, and a classification of existing works according to that system. This paper also discusses the multiple networking standards to support IoT platforms. Zanella *et al.* [8] provided an extensive discussion on network architectures focusing also on transport and application protocols, with practical insights from a concrete deployment. IoT also creates challenges on data management [9] and security [10].

Selected smart cities deployments with IoT solutions are summarized in Table I. City of Things [11], in Brussels, Belgium, encompasses a multitude of sensors (e.g., vehicular environmental stations and parking sensors), whose data is collected via multitechnology gateways or a LoRaWan network. SmartSantander [12] is composed by a large-scale deployment of sensor nodes complemented with citizen participation via mobile devices. Oulu Smart City [13] is a ubiquitous computing testbed that incorporates different wireless networks, large public displays, middleware resources, and monitoring tools.

The distinctive features of the PortoLivingLab platform are as follows. 1) The wide range of targeted urban life aspects, namely environmental, mobility and traffic, mood of citizens, among others. 2) The variety of sensors (personal, vehicle, and static). 3) The scale of the deployment: PortoLivingLab covers one of the largest areas reported to date. 4) Implementing both sensorization and pervasive connectivity whereas most IoT deployments focus on the later, e.g., [11] and [13]). 5) Not requiring dedicated communications infrastructure like [12].

TABLE I
CHARACTERIZATION OF IoT PLATFORMS IN LITERATURE

	CityOfThings [11]	SmartSantander [12]	OuluSmartCity [13]
Information	Traffic flow, air quality, parking spots	Environment (light, temp., CO), parking spots, citizens perspective	Network and protocol performance
Sensors	Static (no. n/a): parking, gateways; vehicular (no. n/a): environmental	Static (no. 790): parking, environ.; vehicular (no. 150): GPS	Sensors/Gateways: WiFi (no. 1200), Bluetooth (no. 12), IEEE 802.15.4 (no. 12)
Collection	Gateways with IEEE 802.11ac 802.15.4, DASH7, BLE and LoRa (no. 100)	IEEE 802.15.4 GWs (no. 100), celullar	

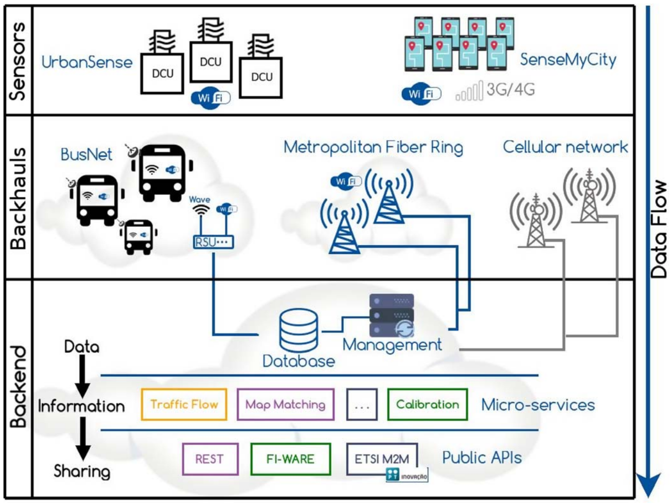

Fig. 1. PortoLivingLab architecture.

III. PORTOLIVINGLAB ARCHITECTURE

PortoLivingLab has been developed under two main premises:

1) Sensing the dynamics of the weather, air quality, noise, and mobility, resorting to data collected by dedicated or opportunistic sensors covering a vast urban area (larger than a district).
2) Building proof-of-concept services and applications that demonstrate the usefulness of the collected data.

The heterogeneity of the sensors and backhauls involved can be structured along the flow of data.

1) *Sensing:* Widespread personal, static and vehicular sensors provided by the urban IoT platforms.
2) *Data Collection:* Heterogeneous communication infrastructure already in place in the city or partly provided by the IoT platforms.
3) *Data Storage, Processing, Sharing, and Hosting Urban Applications/Services:* Backend cloud server.

Fig. 1 depicts the data flow, networking, and physical architectures of PortoLivingLab. Next, we describe the platform following this rationale.

A. Sensing

The platform leverages on static sensors for gathering environmental data (UrbanSense), smartphones for collecting citizens' sensor-enriched trajectory data (SMC), and on-board devices to gather vehicle trajectory and WiFi session data (BusNet). The SMC library turns smartphones into mobile crowdsensing probes, capable of sensing the physical world from the embedded sensors (e.g., location and motion of users and wireless networks) and collect the citizen feedback (e.g., current mood or used modes of transport) through questionnaires. The UrbanSense platform is composed of static data collection units (DCUs) that monitor a number of environmental parameters relating to weather (e.g., temperature, wind speed, and humidity), air quality (e.g., particle and hazardous gases), and quality of life (e.g., noise). The BusNet network is a deployment of OBUs in 600 fleet vehicles (buses, garbage trucks, and street sweepers), capable of collecting the vehicles' mobility traces and quantify connections to the on-board WiFi AP.

B. Data Collection

Urban areas usually feature complex communication infrastructures, such as metropolitan fiber rings, WiFi, and cellular networks, which allow interconnecting the different platform components. In our case, we also resort to a vehicular delay tolerant network (DTN). The sensors of PortoLivingLab reach these backhaul nodes via WiFi gateways. Data collected by UrbanSense is transmitted to the backend server via WiFi APs connected to the fiber ring, municipality WiFi, or vehicular DTN. Data from the SMC probes (smartphones) is transmitted opportunistically on WiFi connectivity (public or private), or via cellular communications at the participant's discretion. The vehicular network BusNet supports data forwarding using V2X communications and DTN, and is also accessed via WiFi.

C. Storage, Processing, Sharing

These tasks are carried out by the backend infrastructure. Incoming raw sensor data is stored in per platform databases. The data models of the different sources are similar to allow easy spatial and/or temporal data fusion. Additionally, the backend server hosts a number of services that provide ready-made data analysis, filtering, and preprocessing (see Section VI). Data sharing with external entities is performed via publish-subscribe middleware for data streams and via a RESTful API for historic datasets. Received sensor or processed data is made publicly available in real-time through standard M2M [14] and FI-WARE [15] interfaces for third-party development of novel urban applications. Crowdsensed data is only provided in an anonymized and aggregated way, e.g., as a stream of average speed values per street or dataset with per mode origin–destination matrices for specific time frames.

IV. URBAN SENSING PLATFORMS

A. SenseMyCity

SMC is a mobile crowdsensing research tool to gather sensor data from participants' smartphones. The SMC sensing probe is a smartphone application that gathers sensor data from the device and from user input. The SMC framework also includes an efficient and secure mechanism of data transmission, described in Section V-B, and ensures privacy and anonymity of collected data, as discussed in Section VI-A. The participants can access, visualize, and download their own data through a Web frontend. A thorough description of SMC is given in [16].

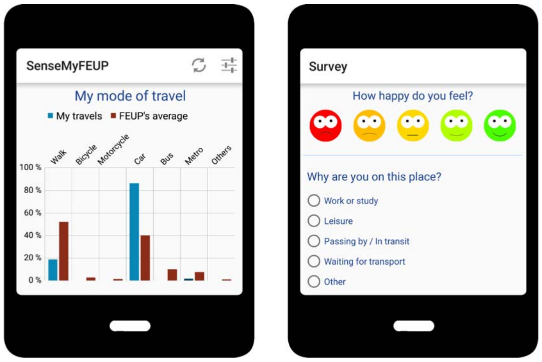

How happy do you feel?

Why are you on this place?

○ Work or study
○ Leisure
○ Passing by / In transit
○ Waiting for transport
○ Other

Fig. 2. Example SMC user interface and mood questionnaire.

1) SMC Mobile Library and Application: The mobile module is a library that enables configuration of the following.

1) *Sensor Sources to Collect:* Data from embedded sensors (e.g., GPS, WiFi, accelerometer, and magnetometer) and external sensors connected via Bluetooth (e.g., on-board diagnostics (OBD) and wearable cardiac sensors).
2) Event triggered questionnaires (e.g., asking travel mode after a trip and randomly asking participants their mood).

User interface and engagement mechanisms (see Fig. 2) are independent and part of the application.

Mobile crowdsensing requires high usability to lower adoption barriers. Usability translates into low battery consumption and limited intrusiveness, which were important driving forces in the design of the SMC mobile library. We measured energy requirements of different sensor configurations in multiple smartphones [16], in order to define a viable set of sensors and their sampling rates. In default settings, location data is collected at 1 Hz and other data, such as nearby WiFi hotspots and user activity, is collected passively, consuming around 5% of battery charge per hour. To reduce the amount of time the application is sensing (and consequently battery consumption), the library automatically starts and stops data gathering based on the smartphones movement and detected activity. Concretely, a displacement higher than 150 m within 3 min triggers a sensing session. Our tests show that the app consumes less than 1% of battery charge per hour when not collecting data. To reduce intrusiveness, the participant only needs to interact with the app for answering questionnaires. Even this is not mandatory, as new questionnaires replace older unanswered ones.

2) Deployment and Challenges: The framework can be configured for different purposes. It has been used in various limited-scope research works, such as studying police officers' and bus drivers' stress [17] or to develop fuel consumption estimation models [18], [19]. We validated that the aforementioned low intrusiveness mechanisms increase the participation, reduce churning, and increase the amount of collected data in three large-scale data collection campaigns. We are currently assessing the data quality for the three campaigns. A major challenge for a crowdsensing tool is participant engagement. We explored three mechanisms in the data collections, and further details can be found in [3].

B. UrbanSense

The UrbanSense platform is a city-wide platform for continuous environmental monitoring. A detailed description can be found in [4]. It is composed of monitoring units which we named DCU deployed in 19 chosen locations within the city.

1) DCUs Hardware and Sensors: A DCU is composed by a power supply, a processing unit (Raspberry Pi), sensors, and a control board to interface the processing unit and the sensors. The whole setup, except for the sensors, is enclosed in a water-proof casing. The sensors are of three classes. 1) *Meteorological:* Thermometer, hygrometer, wind vane, anemometer, rain gauge, lux meter, and solar radiation. 2) *Quality of Life:* Sound level meter. 3) *Air Quality:* Particulate matter sampler, CO, NO_2, and O_3 gaseous meters. The wind vane, anemometer, and rain gauge are mounted in a dedicated structure for freedom of movement. The sound level meter is mounted below the main casing, to ensure exposure to environmental noise and protection from rain. The remaining sensors are enclosed in a separate vented shelter (to allow air flow) mounted on top of the main casing. The DCU also contains two WiFi USB adapters: one for permanent connection to nearby infrastructural or cellular APs used mainly for management although it could also be used for data; and a second one for opportunistic connection to the on-board APs of the BusNet platform. The exterior of a DCU and its breakdown into composing elements is shown in Fig. 3.

2) Deployment and Challenges: A set of tentative locations for deployment were provided by environment and acoustics experts as representative of typical urban settings relevant for environmental characterization: high traffic, residential, waterside, park, and touristic. Deployment started in August 2014 and one year later we had 23 deployed units. Currently, 19 DCUs are still functioning. Full-scale permanent operation of the platform is challenging due to three main reasons.

1) *Dependence on Third-Party Infrastructure and Equipment:* All communication infrastructures between DCUs and the backend server is kept by partner institutions. Failures are disruptive and the recovery time is highly variable.
2) *Unforeseen Limitations of Design Options:* DCUs use a Raspberry Pi for processing, chosen due to the availability of a wide know-how base and support community, but we found it to fail often due to SD card wear down.
3) *Equipment Wear Due to Weather Conditions:* Two units installed near the ocean experienced accelerated corrosion compared with units elsewhere in the city, and the others are bound to follow anytime now.

C. BusNet

The BusNet infrastructure is a large-scale deployment of OBUs in vehicles and infrastructural road-side units (RSUs) that enables vehicle-to-vehicle and vehicle-to-infrastructure communication, in Porto, Portugal. Currently, more than 600 fleet vehicles are equipped with OBUs, pertaining to the public transportation authority (more than 400 buses) and waste disposal department (garbage-collection and road-cleaning trucks). Additionally, more than 50 RSUs have been deployed. The GPS traces and metadata of passenger WiFi connection are collected by each OBU and stored in the backend infrastructure. The network is operated by the PortoLivingLab project and Veniam, a spin-off from the Universities of Aveiro and Porto and the Instituto de Telecomunicações. The technology was published in [20].

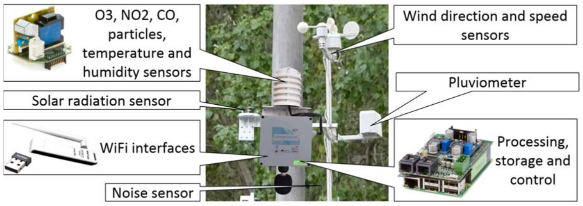

Fig. 3. UrbanSense DCU.

1) Hardware and Technologies: The OBUs are equipped with a processing unit, GPS antenna, cellular module, and wireless interfaces for IEEE 802.11b/g/n (WiFi) and 802.11p (DSRC). BusNet nodes connect to each other via DSRC, and to Internet via RSUs or cellular links. The handover management between similar and different technologies follows the strategy described in [21]. WiFi is also available at BusNet nodes so that bus passengers and external users can connect to the OBU. The RSUs are equipped with a wireless DSRC interface and wired connection to the Internet access. The wired connection uses the technologies available on-site, of which we highlight the municipality-owned fiber ring. Fig. 4 shows the hardware.

V. COLLECTION BACKHAULS AND PROTOCOLS

A. Network Architecture and Backhauls

The network architecture of PortoLivingLab is shown in Fig. 1. PortoLivingLab leverages communication infrastructures already in place in the city and managed by third parties, namely the metropolitan ring fiber (managed by the municipality and accessible via public outdoor wireless APs), the BusNet communication services (operated by Veniam), and a cellular network. All sensing platforms reach the backhauls via wireless links. The backend cloud server features dedicated connections to the various backhauls (Internet and metropolitan fiber ring, that also reaches the BusNet RSUs).

Due to their different nature, the three sensing platforms use distinct backhauls. The SMC mobile application is given access to the Internet (and ultimately to the backend server) at the user's discretion, via a cellular connection or connection to a WiFi network. The UrbanSense platform mainly uses two backhauls: the municipality-owned fiber ring via static WiFi APs and the DTN service provided by BusNet [22]. Cellular hotspots are used sporadically, where the municipality WiFi is not available. BusNet is itself primarily a communication backhaul, as OBUs feature cellular and IEEE

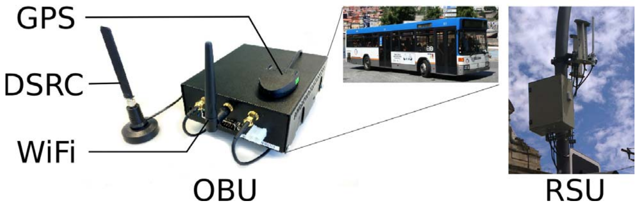

Fig. 4. BusNet platform hardware.

802.11p technologies for Internet access, available via the mobile WiFi hotspots. The latter technology also supports mesh and multihop communication, on top of which a DTN service is provided. This service supports transport of nonurgent data from clients to RSUs, and ultimately the Internet, at low cost (since the cellular connection is not used) but with high latencies, as can be seen in [4]. In addition, it only supports broadcast in the downlink, i.e., unicast addressing from server to clients is not available.

B. End-to-End Protocols

As the communication backhauls in PortoLivingLab are operated by third parties, they are handled as black boxes and without the assumption of continuous connectivity. Two design options of PortoLivingLab protocols and data management are motivated by this setup. 1) Sensors interact directly with the backend (i.e., the backhaul operates as a gateway to the sensors and not as a broker). 2) Sensors store data locally and synchronize it with backend storage whenever there is connectivity.

The SMC mobile application has two options to reliably synchronize local data with the backoffice storage: either continuously, using whatever connectivity is available (including cellular), or opportunistically on WiFi connections only. To support sporadic communication, we developed an asynchronous protocol with end-to-end reliability that implements authentication and encryption for added security. The authentication phase is performed in a single encrypted packet and response (see [16]). During the data transmission phase, the sensor data is serialized, compressed, and then encrypted with an exchanged key. Application layer acknowledgments inform the library that the data has been persistently inserted into the server database and local data can be safely deleted. For efficient use of opportunistic connectivity, several packets can be in transit simultaneously.

The data transfers in UrbanSense are handled differently depending on the used backhaul. The metropolitan fiber ring, accessible to the UrbanSense DCUs via outdoor APs, supports bidirectional communication, and thus an approach similar to SMC is used. Data sent by DCUs is registered locally and not deleted until acknowledged by the backend. At the backend, a service receives the data and issues the acknowledgments. The BusNet DTN service only supports broadcast in the downlink and, as such, end-to-end ACKs cannot be delivered to the DCUs. In turn, a *custody transfer* approach is used. OBUs host a service to receive data in CoAP format via WiFi from the DCUs. The acknowledgment is implicit in the ACK sent by the OBU receiving the data. Nevertheless, end-to-end reliability cannot be guaranteed.

VI. DATA MANAGEMENT, PROCESSING, AND SHARING

A. Data Management

A differentiating aspect of our living laboratory is the straightforward fusion of data from the various platforms. Data fusion is facilitated by a spatio-temporal data model common to all platforms. Concretely, there is a column for timestamps in every sensor table of all databases with common data type that also serves as an index. Additionally, we use a similar coordinates representation for location data. Using common data formats allows for fast relational queries between data of different sensing platforms.

The data collected by the SMC platform is of personal nature and thus it is particularly sensitive. We incorporated privacy and anonymity by design in SMC, following the European Union Data Protection Guidelines. Participants in SMC collection campaigns are volunteers, can leave the data collection at any time, and can delete their own data without intermediation. The participants acknowledge and consent the data collection and processing for research purposes in an informed consent. For awareness, participants see a widget in the smart phone when data is being collected. To guarantee data ownership, the frontend supports data management by the participant, including the possibility to download and delete her own data. Anonymity is provided by using third-party authentication, which avoids the need to store usernames and passwords. We only store the participants' hashed email address to support the aforementioned data ownership for a period of at most one year. We also keep a strong separation between raw datasets and processing datasets. The raw data is accessible only to the database administrator. In the datasets disclosed for processing, we use pseudo-anonymization via a random daily user ID, and extract minimal necessary data fields for the task at hand. When session identifiers are necessary, the timestamps are obfuscated.

B. Data Processing Services

Several data processing services were developed and deployed in the backend server for automated raw data processing. Processed data is stored in dedicated database schemas.

1) Fuel Consumption: We developed a fuel consumption estimation service capable of estimating the instantaneous fuel consumption using only GPS data: vehicle speed, acceleration, and road gradient. We used SMC's capability to gather data from external sensors, in this case from an OBD 2 device. We gathered a dataset of vehicular information, such as speed, fuel consumption and other driving metrics, together with location information from the smartphone's GPS. The synchronized data was used to train a model capable of estimating the fuel consumption and carbon footprint from arbitrary GPS traces. Further details can be found in [18] and [19]. This service is used to pinpoint areas of the city with highest aggregate fuel consumption, information that may be useful to urban planners.

2) Map Matching: For some services, like routing, it is necessary to match the collected data to a graph representation of the road network, allowing for processing and correlation of data per street. We map every location data received to the road network provided by the collaborative OpenStreetMap project. For itinerary traces, e.g., collected by SMC participants and BusNet buses, we use an external service that improves accuracy of the matched road positions by matching and validating the whole trace and not point-by-point (https://mapmatching.3scale.net).

3) UrbanSense Sensor Calibration: Sensors degrade with time, and low cost sensors like the ones we used, even more so. To address this issue, we developed an innovative strategy that automates centralized and periodic calibration of the deployed sensors. This sensor calibration service is performed entirely at our backend infrastructure, and obviates any changes in the deployed DCUs hardware or software. In essence, data from deployed sensors is compared with data from a movable reference sensor. The movable reference sensor is a regular DCU that is regularly calibrated against high precision reference sensors (that are costly and unpractical to carry). By collecting data from both deployed and reference DCUs at the site, a model is obtained via fitting and stored as a calibration configuration. Distinct calibration configurations are stored per sensor, per DCU, and per operation period. Sensor data is calibrated at the backend using the stored configuration, generating a new data stream and a new table. We keep the original parameters to account for possible software errors.

C. Data Sharing

The gathered data is made available via three different APIs.

1) *Standard RESTful APIs for Request-Response Data and Services:* The BusNet and the SMC anonymized datasets are available through RESTful GET methods.
2) *ETSI M2M Publish-Subscriber [14] (Currently Branded As OneM2M):* The sensor data flows from the UrbanSense testbed are exposed through an ETSI M2M broker (NSCL), which can be accessed using a M2M network application. All pub-sub flows can also be queried with GET methods.
3) *Orion Context Broker:* Implementation of the FIWARE publish/subscribe context broker generic enabler, using FIWARE NGSI (Next Gen. Service Interface) API v2.

The use of standard IoT data access and sharing methods is essential for interoperability. In [23], a strategy for machine-to-machine IoT interoperability is proposed.

The service provider and operator of the BusNet vehicular network provides a RESTful API for accessing the BusNet dataflows that can be accessed with a Java Web Token (JWT, RFC 7519) obtained from Veniam. Data is available for buses, and garbage-collection and road-cleaning trucks, regarding real-time vehicle trajectory data and associations to on-board access point, among other information.

As for SMC, three types of historical data are available. 1) SMC participants' trip data can be accessed through APIs only by the participant after authentication. 2) SMC geo-clustered and k-anonymized data ($k = 3$), such as the city current mood map, travel mode statistics, or average speed. 3) Other SMC anonymized datasets are also available, with no links between users, trips, and time of day, such as speeds per road for a certain period of time, and detected WiFi hotspots' locations.

TABLE II
PLATFORM SUMMARY AND DATABASE CHARACTERIZATION

	SenseMyCity	UrbanSense	BusNet
Type	Crowdsensor	Static sensors	Vehicular nodes
Launch year	2011	2015	2014
No. nodes	677 participants	19 DCUs	608 OBUs/RSUs
No. entries	370M	68M	1M/day (≈700M)
Information	Citizens mobility and perspective	Environment	Passenger flows and connectivity
Raw data	Location, motion, WiFi hotspots, questionnaires, external (e.g. OBD)	Temperature, humidity, wind, rain, solar radiation, luminosity, noise, particles, NO_2, O_3, CO	Bus location, passengers WiFi connections

All UrbanSense raw sensor data flows to the above mentioned brokers. Aggregated data, by default hourly and daily, and alerts based on a given set of thresholds for a sensor or a group of sensors are planned, but are not yet available. Subscription to the respective resource triggers the reception of sensor values as they are sent by the DCU. Authentication on the Orion Context Broker is based on FIWARE Laboratory account token and, in case of NSCL, a certificate must be requested.

VII. BENCHMARK AND ILLUSTRATIVE USE CASES

The scope and dimension of the PortoLivingLab platform is summarized in Table II, enabling comparison with other platforms using Table I. PortoLivingLab features a dimension comparable to the largest reported platform (see [12]) in the number of active devices, although it is the only one providing 608 connected mobile sensing probes, and covering the full scope of information collected by the set of the IoT platforms presented in Section II. We are unable to compare the dimension of the datasets of the different platforms since these have not been disclosed by other works.

We now showcase proof-of-concept use cases that demonstrate the usefulness of data generated by PortoLivingLab for smart city studies. We present two cases per platform—SMC, UrbanSense, and BusNet—and a final use case leveraging data from BusNet and UrbanSense. An outlook of future lines of research with the available datasets concludes the section.

A. Citizens Mood Map

In 2015, the SMC platform was configured for the SenseMyMood data gathering campaign. It consisted of an exploratory study of the connection between places, times of day, and emotions, and served as engagement mechanism at the same time. The SenseMyMood app automatically gathered mobility data during the user's trips, and sporadically requested participants to self-assess their happiness and other emotions. A notification for a new questionnaire was shown randomly up to twice per day, asking users their current activity—work or study, leisure or resting, in transit, waiting for transport, other—and to report their happiness in a scale of 1 (very sad) to 5 (very happy). Optionally, users could also rate their self-assessment of the six basic emotions of the Ekman model—joy, sadness, fear, surprise, disgust, and anger. A snapshot of the smartphone sensors was collected when answering

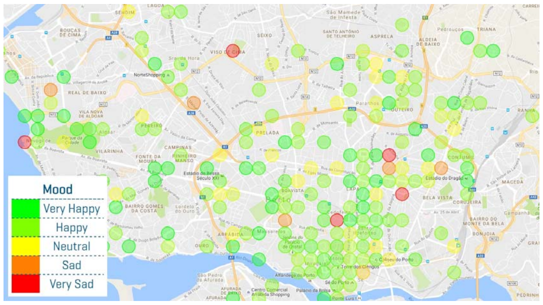

Fig. 5. SenseMyMood: average happiness reported by participants.

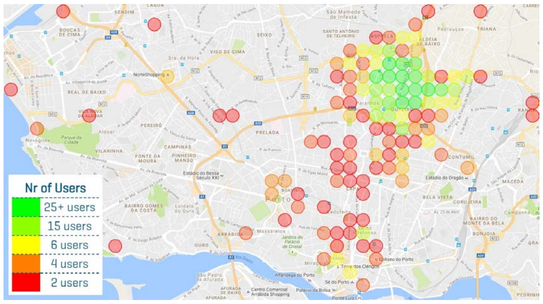

Fig. 6. SenseMyFEUP: origin/destination of trips going to/from FEUP.

these mood reports, providing us with the instantaneous location and speed, ambient noise level, nearby WiFi hotspots, and other sensors data if available. Around 3100 responses in Porto were used to create the clustered citizens mood map shown in Fig. 5, and allowed us to correlate the user's activities with the self-assessed emotions. As an example, we were able to clearly associate lower happiness levels with traffic and waiting for transport, which may support traffic management decisions. A similar application can be found in [24], although with much inferior geographical density of users, and thus not enabling an analysis in a city scale. Traditionally, such studies were done via resource-consuming questionnaires, reaching much smaller samples and coarser spatial and temporal granularity.

B. Digital Mobility Survey

In 2016, the SenseMyFEUP data gathering campaign was carried out in the University of Porto campus using the SMC infrastructure. The mobile application gathered mobility data automatically during a one-month period (April), and participants were asked to report which transportation modes they used after each trip. The campaign dissemination was done on the engineering school campus with a regular population of 9000 people, via flyers, posters, and emails targeting students, academics, and other staffs. Engagement was improved by randomly drawing nonmonetary weekly prizes among participants. The SenseMyFEUP mobile application was adopted by 239 members of the faculty community who contributed with 8700 responses of used transportation modes during that period. We obtained information about the community's transportation modes, duration and distance to and from the school, and most common origins and destinations, as shown in Fig. 6. Traditionally, such mobility studies are done using mobility surveys repeated in time scales of decades. Recent works use vehicle counters [25], cellular network datasets [26], or taxi datasets [27]. Crowdsensing enables finer spatial and temporal granularity and is not biased by particular transportation modes, but in our campaigns we observed a sample bias toward young adults, i.e., university students. The results are being compared with traditional surveys to evaluate its accuracy on travel demand assessment. The gathered transportation mode survey data was used as ground truth to train a machine learning algorithm to automatically detect transportation mode from standard location traces [28].

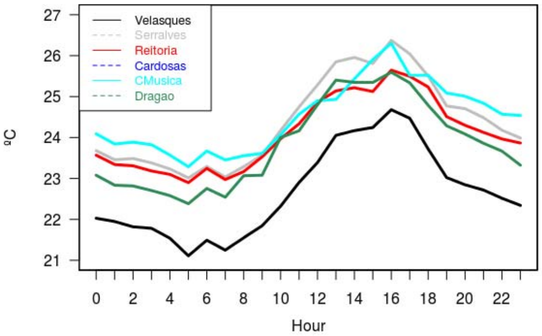

Fig. 7. Average temperature per hour of day during 18 days.

C. Differentiated City-Wide Climate Analysis

From the weather information collected so far, we are able to create a spatial characterization of climate and air quality in the city of Porto. As an example, we evaluated the difference in temperatures at different city locations, during the same time period. We retrieved the data collected by the UrbanSense units at monitored areas for the period of October 2, 2015 to October 19, 2015, and produced the hourly average temperature from calibrated data. In Fig. 7, we observe that all locations exhibit similar temperature profiles over the day, and that one location (Velasquez) presents a consistent difference of around 1 °C to the remaining locations. Similar studies were carried out by IoT platforms reviewed in Section II, namely [11] and [12]. Traditionally, weather is observed on very few locations by very high quality and proportionally costly sensors. Smart city-IoT promises to complement traditional physical weather models with data-based local models for finer granularity.

D. Impact of Forest Fires in Air Quality

On the week of August 7, 2016 to August 13, 2016, a number of large forest fires surrounded the city of Porto, and the air was noticeably filled with smoke. We searched the UrbanSense database for air quality-related parameters at week-long time intervals before, during, and after these dates, and compared the evolution of those parameters to validate this event. Fig. 8 shows the concentration of particulates, ozone, luminosity, and solar radiation from an UrbanSense unit during, before, and after the event. We observe that both ozone concentrations and solar intensity (directly linked to ozone formation rate) are lower during the fire period. The average luminosity also decreased, possibly due to the airborne particles. We note an increase in the concentration of particles of different sizes during that particular week, which in turn can be a consequence of the fire smoke that covered the city. This example demonstrates

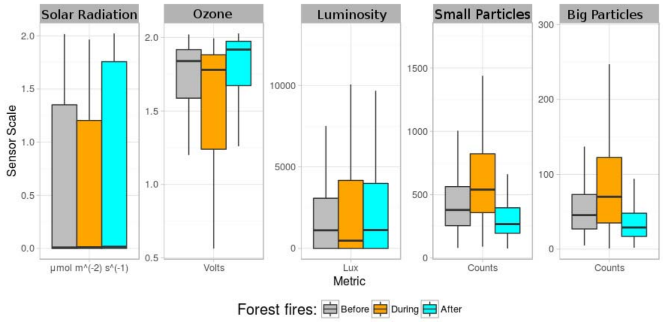

Fig. 8. Environmental indicators before, during, and after the fire period.

that urban sensing platforms can identify periods susceptible of triggering different traffic management measures to keep ozone or particle levels within acceptable limits for human health. Air quality is usually monitored and modeled similarly to the weather, and for large time scales, e.g., yearly reports. Low cost sensors, although incurring in lower accuracy, enable finer grained observations allowing monitoring and modeling for much more locations.

E. Passenger Flows From WiFi Connections

The metadata of the WiFi connections of passengers to the on-board APs of BusNet, alongside with the GPS traces recorded by the buses' OBUs, can be used to construct the origin–destination matrices of passengers. Understanding the flows of passengers in public buses helps to improve the service, e.g., by identifying routes that require an increase in bus frequency. By matching the start and end time of the WiFi connections with the timestamps of the OBU position traces available in the PortoLivingLab datasets, and using visualization tools developed for this purpose, we created an animation[2] of passenger flows in the city of Porto for one week in October 2014. Fig. 9 presents a still frame of that video. Similar work can be found in [29], although at a smaller scale.

F. Mobility Patterns of Buses

The speed of buses can also be analyzed to extract information about the traffic flows in the city. This information can be used, e.g., to find spots where speed is chronically low. The average speed of buses, at road segments in the city of Porto, was computed from the BusNet datasets and, using a threshold rule, the locations with lower average speed (critical spots) were identified. Fig. 10 presents the average bus speed and critical spots in Porto during the period 7–11 A.M. of June 9, 2014. Using mobility traces for traffic estimation has been explored in [30], as well as the extrapolation of global conclusions from a subset of nodes (e.g., buses) [31].

G. Performance of Delay Tolerant Networking

The BusNet serves as a communication backbone to the UrbanSense platform, as DCUs offload their sensor data to OBUs. This infrastructure-to-vehicle (I2V) collection mechanism is opportunistic and circumstantial by nature: opportunities may vary from site to site due to different bus frequencies, number of nearby bus stops, etc. To learn how the performance of I2V collection differs between locations, we recorded at several DCU sites the duration of contacts between DCU and OBUs and the number of packets exchanged in each contact. This data can be further processed to get the average values of contact duration and number of contacts per day, and number of exchanged bundles per connection over week-long or month-long periods of time. This will tell us, e.g., if a location served by many connections with short contact time outperforms another with few connections and longer contact times. Fig. 11 presents the ratio of the above metrics with respect to the platform-wide maximum values, at several DCU locations and for the period between October 2, 2015 and October 22, 2015. Our vehicular DTN platform also allows to study how vehicles can cooperate to gather the data toward the infrastructure reliably and with small delay [32], and to characterize WiFi performance in I2V links to urban fleet vehicles [33]. Implementations of solutions in urban environments can be traced back to [34]. We did not find works reporting large-scale vehicular delay-tolerant collection from mobile or static road-side nodes, to the best of our knowledge.

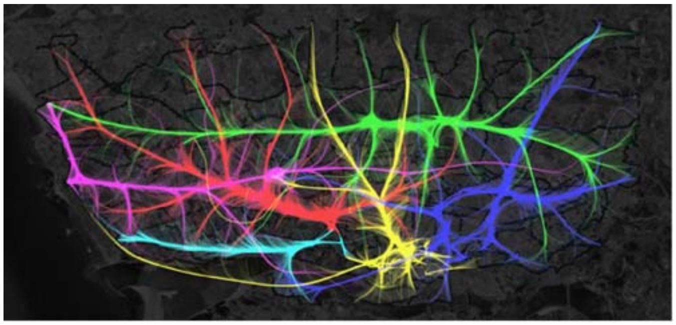

Fig. 9. Passenger flows from WiFi connections.

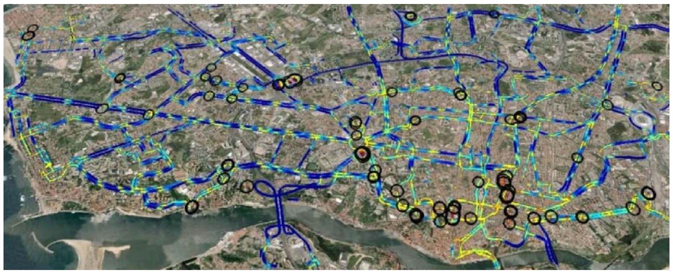

Fig. 10. Average bus speed. Color scale indicates average speed: blue for high and red for low. Black circles are critical locations.

H. Future Directions for Data Analysis

Next steps will focus on the combined analysis of the PortoLivingLab dataset. The three platforms contain partially overlapping or complementary information that, after analysis, can validate some initial hypothesis. The test hypotheses can come from empirical observations or previous studies. For example, if one considers a public event, such as a football

---

[2] Full video. [Online]. Available: https://vimeo.com/125892304.

Fig. 11. Performance of I2V collection of UrbanSense DCUs data at locations served by BusNet nodes. Per location, the number of OBU-DCU contacts (left), total contact time (top), and number of exchanged bundles (right) are shown as a ratio to the network's maximum values (black circles).

game or a street concert, we expect to find in that area a large aggregation of people, increased traffic in roadways and public transports, and higher noise levels. From the PortoLivingLab perspective, SMC can provide information about density of people, means of transportation used and even trends in their feelings, BusNet can detect disruption in normal traffic patterns, and noise levels can be measured by UrbanSense sound level meters. In a second example, we observe empirically that when it rains traffic moves slower, jams are more frequent, people are likely to stay indoors, and their mood is more somber. Upon identification of rain periods by UrbanSense rain gauges, we can characterize traffic conditions as well as the people mobility and mood from BusNet and SMC data. We envision the conclusions obtained by the PortoLivingLab platform to provide relevant information to improve the city, not only in a city management level, but also in a citizen level.

VIII. CONCLUSION

We presented PortoLivingLab, three complementary sensing platforms and a common backend infrastructure for data storage and processing, deployed in the city of Porto. We describe the scope and operation of the whole platform, including its architecture combining different types of sensors and the strategies for data collection of diverse information sources, and present the valuable dataset resulting of its operation over a period of almost two years. Deeper understanding of perceived phenomena has been obtained through careful analysis of collected data. Future work will focus on refining and improving phenomenon comprehension and providing the relevant conclusions to the decision makers.

ACKNOWLEDGMENT

The authors would like to thank T. Calçada, C. Pérez-Penichet, and B. Fernandes for their technical contribution to the UrbanSense platform, the team at VENIAM for maintaining BusNet, D. Moura for the visualizations of BusNet data, and V. Ribeiro for his contribution to the SenseMyCity framework. The authors would also like to thank the municipality of Porto and Associação Porto Digital for support in all platforms.

REFERENCES

[1] C. Harrison and I. A. Donnelly, "A theory of smart cities," in *Proc. 55th Annu. Meeting ISSS*, vol. 55, 2011, pp. 1–15.

[2] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Comput. Netw.*, vol. 54, no. 15, pp. 2787–2805, 2010.

[3] J. G. P. Rodrigues, A. Aguiar, and C. Queirós, "Opportunistic mobile crowdsensing for gathering mobility information: Lessons learned," in *Proc. IEEE Int. Conf. Intell. Transport. Syst.*, Rio de Janeiro, Brazil, 2016, pp. 1654–1660.

[4] Y. Luis *et al.*, "Urbansense: An urban-scale sensing platform for the Internet of Things," in *Proc. IEEE Int. Smart Cities Conf.*, Trento, Italy, 2016, pp. 1–6.

[5] C. Ameixieira *et al.*, "Harbornet: A real-world testbed for vehicular networks," *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 108–114, Sep. 2014.

[6] J. Jin, J. Gubbi, S. Marusic, and M. Palaniswami, "An information framework for creating a smart city through Internet of Things," *IEEE Internet Things J.*, vol. 1, no. 2, pp. 112–121, Apr. 2014.

[7] C. Perera, Y. Qin, J. C. Estrella, S. Reiff-Marganiec, and A. V. Vasilakos, "Fog computing for sustainable smart cities: A survey," *ACM Comput. Surv.*, vol. 50, no. 3, pp. 1–43, Jun. 2017.

[8] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, Feb. 2014.

[9] Y. Qin *et al.*, "When things matter: A survey on data-centric Internet of Things," *J. Netw. Comput. Appl.*, vol. 64, pp. 137–153, Apr. 2016.

[10] J. Zhou, Z. Cao, X. Dong, and A. V. Vasilakos, "Security and privacy for cloud-based IoT: Challenges," *IEEE Commun. Mag.*, vol. 55, no. 1, pp. 26–33, Jan. 2017.

[11] S. Latre *et al.*, "City of things: An integrated and multi-technology testbed for IoT smart city experiments," in *Proc. IEEE Int. Smart Cities Conf. (ISC2)*, Trento, Italy, 2016, pp. 1–8.

[12] L. Sanchez *et al.*, "Smartsantander: IoT experimentation over a smart city testbed," *Comput. Netw.*, vol. 61, pp. 217–238, Mar. 2014.

[13] T. Ojala, "Open urban testbed for ubiquitous computing," in *Proc. Int. Conf. Commun. Mobile Comput. (CMC)*, vol. 1. Shenzhen, China, 2010, pp. 442–447.

[14] ETSI. (2016). *ETSI-Internet of Things Website*. [Online]. Available: http://www.etsi.org/technologies-clusters/technologies/internet-of-things

[15] Fiware Community. (2016). *Fiware*. [Online]. Available: http://www.fiware.org/

[16] J. G. P. Rodrigues, A. Aguiar, and J. Barros, "SenseMyCity: Crowdsourcing an urban sensor," *arXiv preprint arXiv:1412.2070*, 2014. [Online]. Available: http://arxiv.org/abs/1412.2070

[17] J. G. Rodrigues, M. Kaiseler, A. Aguiar, J. P. S. Cunha, and J. Barros, "A mobile sensing approach to stress detection and memory activation for public bus drivers," *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 6, pp. 3294–3303, Dec. 2015.

[18] V. Ribeiro, J. Rodrigues, and A. Aguiar, "Mining geographic data for fuel consumption estimation," in *Proc. 16th Int. IEEE Conf. Intell. Transp. Syst.*, 2013, pp. 124–129.

[19] A. Vilaça, A. Aguiar, and C. Soares, "Estimating fuel consumption from GPS data," in *Proc. Iberian Conf. Pattern Recognit. Image Anal.*, Santiago de Compostela, Spain, 2015, pp. 672–682.

[20] C. Ameixieira *et al.*, "An IEEE 802.11p/WAVE implementation with synchronous channel switching for seamless dual-channel access," in *Proc. IEEE Veh. Netw. Conf.*, 2011, pp. 214–221.

[21] D. Lopes and S. Sargento, "Network mobility for vehicular networks," in *Proc. IEEE Symp. Comput. Commun. (ISCC)*, 2014, pp. 1–7.

[22] P. M. Santos *et al.*, "Demo: Platform for collecting data from urban sensors using vehicular networking," in *Proc. 21st Annu. Int. Conf. Mobile Comput. Netw.*, 2015, pp. 167–169.

[23] C. Pereira *et al.*, "IoT interoperability for actuating applications through standardised M2M communications," in *Proc. IEEE 17th Symp. World Wireless Mobile Multimedia Netw. (WoWMoM)*, Coimbra, Portugal, 2016, pp. 1–6.

[24] G. MacKerron and S. Mourato. (2011). *Mappiness*. [Online]. Available: http://www.mappiness.org.uk/

[25] H. Shaoa, W. H. K. Lam, A. Sumalee, and M. L. Hazeltone, "Estimation of mean and covariance of stochastic multi-class OD demands from classified traffic counts," *Transport. Res. Part C Emerg. Technol.*, vol. 59, pp. 92–110, Oct. 2015.

[26] M. S. Iqbal, C. F. Choudhury, P. Wang, and M. C. Gonzalez, "Development of origin–destination matrices using mobile phone call data," *Trans. Res. Part C Emerg. Technol.*, vol. 40, pp. 63–74, Mar. 2014.

[27] L. Moreira-Matias, J. Gama, M. Ferreira, J. Mendes-Moreira, and L. Damas, "Time-evolving O-D matrix estimation using high-speed GPS data streams," *Expert Syst. Appl.*, vol. 44, pp. 275–288, Feb. 2016.

[28] J. G. P. Rodrigues, J. P. Pereira, and A. Aguiar, "Impact of crowdsourced data quality on travel pattern estimation," in *Proc. 1st ACM Workshop Mobile Crowdsens. Syst. Appl. (CrowdSenSys)*, Delft, The Netherlands, 2017, pp. 38–43.

[29] V. Kostakos, T. Camacho, and C. Mantero, "Wireless detection of end-to-end passenger trips on public transport buses," in *Proc. 13th IEEE Conf. Intell. Transp. Syst.*, 2010, pp. 1795–1800.

[30] J. C. Herrera *et al.*, "Evaluation of traffic data obtained via gps-enabled mobile phones: The mobile century field experiment," *Transp. Res. Part C Emerg. Technol.*, vol. 18, no. 4, pp. 568–583, 2010.

[31] R. Bertini and S. Tantiyanugulchai, "Transit buses as traffic probes: Use of geolocation data for empirical evaluation," *J. Transport. Res. Board*, vol. 1870, no. 1, pp. 35–45, 2004.

[32] M. A. Khan, S. Sargento, and M. Luis, "Data collection from smart-city sensors through large-scale urban vehicular networks," in *Proc. IEEE Veh. Technol. Conf. Fall*, Toronto, ON, Canada, Sep. 2017.

[33] P. M. Santos, T. Calçada, S. Sargento, A. Aguiar, and J. Barros, "Experimental characterization of I2V Wi-Fi connections in an urban testbed," in *Proceedings of the 10th ACM MobiCom Workshop on Challenged Networks.* New York, NY, USA: ACM, 2015, pp. 5–8.

[34] B. Hull *et al.*, "CarTel: A distributed mobile sensor computing system," in *Proc. 4th ACM SenSys*, Boulder, CO, USA, Nov. 2006, pp. 125–138.

Pedro M. Santos received the M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of Porto, Porto, Portugal, in 2009 and 2017, respectively.

He is a Post-Doctoral Research Fellow and Invited Assistant Professor with the University of Porto. He is currently a Researcher and an Invited Assistant Professor with the University of Porto. His current research interests include physical layer, vehicular networking, mobile computing, and computer architectures.

João G. P. Rodrigues received the M.Sc. degree in electrical and computer engineering from the University of Porto, Porto, Portugal, in 2009, where he is currently pursuing the Ph.D. degree.

He is with the Instituto de Telecomunicações, Porto, researching data gathering and mining in intelligent transportation systems. His current research interests include sensor networks and ITS.

Susana B. Cruz received the M.Sc. degree in electrical and computer engineering from the Faculty of Engineering, University of Porto, Porto, Portugal, in 2010, where she is currently pursuing the Ph.D. degree.

She is with the Instituto de Telecomunicações, Porto. Her current research interests include wireless sensor networks and intelligent transportation systems, vehicular networks, localization, and cooperative approaches.

Tiago Lourenço received the Licenciatura degree in computer and telecommunications engineering from the Technology and Management School, Lamego Polytechnic Institute, Lamego, Portugal, in 2009, and the M.Sc. degree in information engineering from the Faculty of Engineering, University of Porto, Porto, Portugal, in 2016.

He was a Researcher with the Instituto de Telecomunicações, Porto. Currently, he works with Ericsson, Lund, Sweden.

Pedro M. d'Orey received the Licenciatura degree in electrical and computer engineering from the University of Porto, Porto, Portugal, in 2004, the M.Sc. degree in telecommunications from Queen Mary University of London, London, U.K., in 2008, and the Ph.D. degree in telecommunications from the University of Porto, in 2014.

He is currently a Researcher with the Instituto de Telecomunicações, Porto, and an Invited Assistant Professor with the University of Porto. His current research interests include vehicular networks, automated vehicles, and data mining.

Yunior Luis received the B.Sc. degree in computer science from the Faculty of Mathematics and Computer Science, University of Havana, Havana, Cuba, in 2014. He is currently pursuing the Ph.D. degree at the University of Porto, Porto, Portugal.

Since 2015, he has been a Researcher with the Faculty of Engineering, University of Porto. He is with the Instituto de Telecomunicações, Porto. His current research interests include machine learning and databases.

Cecília Rocha received the graduation degree in civil engineering from the Faculty of Engineering, University of Porto (FEUP), Porto, Portugal, in 1994, and the M.Sc. degree in buildings construction and Ph.D. degree in civil engineering from the University of Porto, in 1996 and 2010, respectively.

She has been an Assistant Professor with FEUP, since 2010, where she is a Researcher with the Centro de Investigação do Território, Transportes e Ambiente. Her research has been published in noise conferences and journals. Her current research interests include acoustics, namely in environmental and building acoustics, and urban sustainability.

Sofia Sousa received the Licenciatura degree in chemical engineering and Ph.D. degree in environment engineering from the University of Porto, Porto, Portugal.

She is a Researcher with the Faculty of Engineering, Laboratório de Engenharia de Processos, Ambiente, Biotecnologia e Energia, University of Porto. Her current research interests include air quality, particularly monitoring, modeling, and its impact on human health.

Sérgio Crisóstomo received the M.Sc. degree in electrical and computer engineering and Ph.D. degree in computer science from the University of Porto, Porto, Portugal, and the Ph.D. degree in technischen wissenschaften from Alpen-Adria-Universitaet Klagenfurt, Klagenfurt, Austria (joint Ph.D. degree).

He is an Assistant Lecturer with the Computer Science Department, School of Sciences, University of Porto, and a Researcher with the Instituto de Telecomunicações, Porto. His current research interests include wireless and sensor networks, on positioning and navigation techniques and on cooperation, self-organization, and information dissemination in complex networks.

Cristina Queirós received the graduation and Ph.D. degrees in psychology from the Faculty of Psychology and Sciences of Education, University of Porto (FPCEUP), Porto, Portugal, in 1988 and 1997, respectively.

She is an Assistant Professor with FPCEUP, where she is a member of the Centro de Psicologia da Universidade do Porto and the Co-Director of the Laboratório de Reabilitação Psicossocial. She contributes to interdisciplinary projects where emotions/stress are measured in first-responders and health professionals. She has authored or co-authored several papers and chapters on emotions and stress. Her current research interests include emotional states related to the environment and its impact on quality of life (e.g., stress elicited by modern lifestyles, trauma, and burnout across rescue workers).

Susana Sargento (S'96–M'07) is an Associate Professor (with "Habilitation") with the University of Aveiro, Aveiro, Portugal, and the Instituto de Telecomunicações, Porto, Portugal, where she leads the Network Architectures and Protocols Group (http://nap.av.it.pt). She has been involved in national, CMU-Portugal, and FP7 projects, and in the EU Coordinated Support Action 2012-316296 "FUTURE-CITIES." She co-founded a spin-off company of the Universities of Aveiro and Porto, Veniam Inc. (www.veniam.com), Mountain View, CA, USA, which builds seamless vehicle-based Internet infrastructure. Her current research interests include self-organized networks, ad-hoc and vehicular network mechanisms and protocols, and content distribution networks.

Ms. Sargento was a recipient of the 2016 EU Prize for Women Innovators.

Ana Aguiar (S'94–M'98) received the graduation degree in electrical and computer engineering from the Faculty of Engineering, University of Porto (FEUP), Porto, Portugal, in 1998, and the Ph.D. degree in telecommunication networks from the Technical University of Berlin, Berlin, Germany, in 2008.

She has been an Assistant Professor with FEUP since 2009. She also contributes to several interdisciplinary projects in the fields of intelligent transportation systems and well being (stress). She began her career as an RF Engineer researching cellular operators. She was with Fraunhofer AICOS, Porto. Her research has been published in several IEEE and ACM conferences and journals. Her current research interests include wireless networking and mobile sensing systems, specifically vehicular networks, crowd sensing, and machine-to-machine communications.

Dr. Aguiar is a Reviewer for several IEEE and ACM conferences and journals.

João Barros (S'98–M'04–SM'11) received the undergraduate degrees in electrical and computer engineering from the University of Porto, Porto, Portugal, and the Universität Karlsruhe, Karlsruhe, Germany, and the Ph.D. degree in electrical engineering and information technology from the Technische Universität München, Munich, Germany.

He is currently a Professor Catedrático (Full Professor) of electrical and computer engineering with the University of Porto. He was the Founding Director of the Instituto de Telecomunicações, Porto, and the Director of the CMU-Portugal Partnership Program from 2009 to 2012. He is the co-founder and the CEO of Veniam Inc., Mountain View, CA, USA, a start-up that specializes in vehicular communication technologies, and the co-founder of Streambolico, Porto, which commercialize wireless video technologies.