NOSE: A NOmadic Scalable Ecosystem for pervasive sensing, computing and communication

Audrey Dupont, Yann Bocchi, and Gianluca Rizzo

HES SO Valais, Switzerland, `name.surname@hevs.ch`

Abstract. The Smart City paradigm is progressively shaping the way we interact with other citizens, with institutions, as well as the way in which all kind of resources are managed in an urban setting. However, the high cost of dedicated sensing, computing and communication infrastructure, represents one of the main obstacles to the adoption of the Smart City paradigm in small cities and large, distributed districts, where resides the vast majority of the world population. In this paper we present a first implementation of a platform for nomadic sensing, which exploits the moving infrastructure of a city in order to implement the sensing, computing and communication functionalities required by smart city services, in a cost-effective way. This paper describes the architecture of the NOSE system, and the main lessons learned during the first steps of its implementation, deployment, and experimental validation. In particular, we show that the main technical challenges are due to the lack of devices conceived for reliable mobile operation (sensing, computing) and to the difficulty in achieving high levels of accuracy in measurements while moving.

1 Introduction

Urban districts have a major impact on the economic and social development of nations. They are genuine platforms where people live, where companies have their business and in which numerous services are provided. Moreover, they are major centers of consumption of resources. Currently, they consume 75% of the worlds resources and energy and generate 80% of greenhouse gases, occupying only 2% of the worlds territory.

From a descriptive point of view, a Smart City is an urban space with infrastructures, and intelligent networks and platforms, with millions of sensors and actuators, among which people themselves and their mobile telephones must also be included. Indeed, in the near future, those districts which are not able to take advantage of the instruments which the Smart City paradigm offers are bound to pay a high price, in an increasingly integrated economical context.

The first requirement for the implementation of the Smart City paradigm is a pervasive infrastructure of sensing, computing and communications, together with instruments for data intelligence. Such a condition is typically satisfied by large urban environments, where efficiency gains justify the CAPEX and OPEX of the infrastructure. Conversely, medium/small cities, towns and districts with a sparse distribution of population (and city districts in developing countries) are presently cut out of the smart city paradigm and of its potential benefits, with a negative impact on their present economy as well as on their potential for economic growth and expansion. Small cities with limited budgets, towns and sparse agglomerations, regions with a very fractioned administrative will suffer even more than they do today from the networking effect, which already brings opportunities to concentrate in a few large urban agglomerations. This may bring to slow economic growth or economic depression, and to accentuate the migration of population, capitals, and opportunities towards big, smart cities.

With hardware prices dropping due to commoditization, the main reason for the high CAPEX and OPEX of a distributed sensing and computing infrastructure are the costs of installation and maintenance. Moreover, traditional designs and deployments are ad-hoc for a single service, making it harder to justify infrastructure investments and hampering large scale adoption of smart urban services.

The starting point of our work is the observation that a large portion of sensing and monitoring tasks which enable smart city services do not require real time operation and reporting of data, nor a high temporal resolution. For the majority of such tasks, a few readings of the parameters a day are enough to extract the information necessary to feed smart city services. The idea at the basis of NOSE is to decrease drastically the amount of required sensing and communication infrastructure through mobile sensing, by taking advantage of the already existing nomadic part of the infrastructure of a city or an urban agglomeration, represented by public buses, and by all services on wheels, such as mail/parcel delivery, taxis, and similar. Indeed, almost all such nomadic infrastructure is nowadays connected to the Internet, for different reasons. And all those connected devices are usually already equipped with a standard set of sensors (GPS, accelerometers, temperature, microphones, cameras, etc). This offers the opportunity to implement a nomadic sensing infrastructure, by using every moving agent as a sensing agent. By coupling every sensor reading with spatial information, moving agents are able to produce context-rich information in the same way as fixed sensing infrastructure is able to.

Differently than typical crowdsourcing solutions (often based on unreliable and occasional contributions from a set of devices which vary over time), such nomadic sensing infrastructure relies on agents moving according to patterns which are well defined and easily predictable, with small variability in time and space. Thanks to its ubiquity and reliability, such nomadic infrastructure constitutes potentially a valid, cost effective replacement (and/or integration, for services with more stringent requirements) for the expensive dedicated fixed sensing computing and communication infrastructure. Finally, by enabling smart services in all those portions of a territory reachable through a vehicle, it allows extending the smart city paradigm cost-effectively well beyond the boundaries of cities, to small towns and up to the most remote locations.

In this paper, we describe a first implementation and evaluation of such nomadic sensing and computing infrastructure. We describe the general architecture, our implementation choices, and the main issues faced in applying such approach.

The paper is organized as follows. In Section 1.1 we review the main results in the state of the art, and in Section 2 we describe a set of reference applications for our platform. In Section 3 we outline the system architecture, and we describe the main implementation choices. Section 4 describe the main lessons learned in implementing and evaluating the platform and the sample services. Finally, Section 5 concludes the paper.

1.1 Previous work

Research has long aimed at developing prototypes and solutions for intelligent cities. MIT Smart Cities Lab [6] focuses upon intelligent, sustainable buildings, mobility systems (GreenWheel Electric Bicycle, Mobility-on-Demand, Citycar, Wheel Robots). The IntelCities [9] research consortium developed solutions for electronic government, planning systems and citizen participation; Urenio [15] has developed a series of intelligent city platforms for the innovation economy. The Smart Cities Academic Network [1] is working on e-governance and e-services in the North Sea region. In Europe, one of the most advanced SC prject is the Smart Santander project [7].

Some of the main technical issues in mobile sensing are data heterogeneity, abundance of data, intermittent network connectivity, accuracy of measurements. Several works have addressed these issues. The most representative example of a platform running on buses, and delivering a set of IT services to a sparse community is given by First Mile Solutions [4]. This company implements systems for DTN services such as email, messaging and even offline web search, for rural areas of developing countries, based on cheap mobile infrastructure running on buses and other vehicles, assisted by some fixed infrastructure at critical locations. [13] proposes a system for monitoring air pollution in a city based on sensors installed on public transportation vehicles. [14] analyses the main technical problems related to the nomadic approach to urban sensing that we also adopt in the present proposal. [10] designs and evaluates *Pothole patrol*, a low-cost system that monitors the quality of the road surface, based on a laser line. All these works explore different dimensions of the above mentioned issues in nomadic sensing. These works give an idea of where are the core technical issues involved in implementing a nomadic smart city infrastructure. However, they propose platforms which are not designed to be scalable and multi-service, but rather ad-hoc for the given sensing, computing or communication service considered, and hence hard to generalize.

Other works proposed platforms for addressing a variety of issues, related to delay tolerant networking and mobile sensing. The Diverse Outdoor Mobile Environment (DOME) project at University of Massachusetts [16] was the longest-running large-scale, highly diverse mobile systems testbed. Such testbed has been operational since 2004 and provided infrastructure for a wide range of mobile computing research. The MIT CarTel project [12] aimed at designing and prototyping a distributed software system that makes it easy to collect, process, deliver, visualize and analyse data from mobile sensors (cars, phones, etc). A common denominator of such infrastructure is the focus on sensing, computing and communication issues rather than on the service(s) relying on the nomadic infrastructure and their QoS.

2 Nomadic sensing applications

For our implementation, we focused on a category of *smart city sensing applications*, by which several environmental parameters are sensed in a distributed fashion, collected and aggregated in order to build a view of the city environment in terms of spatio-temporal distributions and patterns of environmental parameters such as pollutants (air, water), pollens, temperature, humidity, and so on.

More specifically, we have considered the following services:

- *Snow Chains Alert*. One of the main issues, on the occurrence of heavy snow storms, is to know when it is optimal to mount snow chains. When mounted unnecessarily, they may damage tyres, get consumed more early, while significantly impacting the duration of the trip (the maximum speed allowed with snow chains is typically around $40Km/h$). Conversely, whenever they should be used but they are not (e.g. due to driver inexperience, or to environmental factors which are hard to detect visually, such as presence of ice on the road surface during night time), the risk of accidents increases.
  The goal of such service is to enable road users to make an informed choice, through a (near) real-time view of the status of the road surface along the planned route. The service should issue an alert, a notification showing the parts of the route where snow chains are likely to be needed, and the maximum driving speed at each point of the route, which takes into account the features of the vehicle (e.g. type of tyres) and of the road (e.g. smoothness of road surface, etc). Among the intended users are all drivers of vehicles, but also road maintenance services, police fire brigades, rescue services, accident prevention, tourist services, insurance companies.
- *Safe driving companion*. A key component of a safe driving behavior is the capability of adapting driving choices (such as speed, acceleration) to the condition of the road, and to the context at large (e.g. traffic conditions, presence of pedestrians, etc.). On the onset of weather conditions such as rain or snow, such adaptation choices are based on the experience of the driver, or on coarse indications by dynamic signalling devices on the road, in some cases, which often are given in terms of maximum speed limit, or in terms of generic advices (”slow down”, ”drive carefully”), only aiming at keeping high the attention level of the drivers, rather than helping them in making good driving choices. In 2015, 8087 road accidents in Switzerland have had as a prime cause a poor management of the speed of the vehicle [2]. Of these, 64.9% have involved vehicles sliding or slipping out of control. A higher awareness on road condition, coupled with a better road maintenance could considerably reduce the occurrence of such accidents and decrease their cost for society.
  Helping achieve such awareness is the main goal of the *Safe driving companion* service. It consists in an application which keeps constantly informed the driver not only on the status of the road, but also on the maximum speed considered safe for the vehicle at a given point of the road, as well as giving an indication on the minimum safety distance at that point. Intended users are all vehicle drivers, but also insurance companies, road maintenance services, and the like.
- *Real-time road safety map* This service is mainly intended for urban/district planners and road maintenance planning and management. The issue it addresses is the current lack of accurate information on how safe a given road segment is. Given the complexity of the problem of safety, which involves not only physical factors (such as road surface status) but also human factors, we believe that compiling a map which associates to every point on a road a given measure of risk of accident cannot be done only on the basis of data such as the history of past accidents, drivers/experts feedback, road shape and characteristics.
  The service consists in a map showing, real-time information, historic data, and forecasts on the estimated grip in every road segment in a given area,

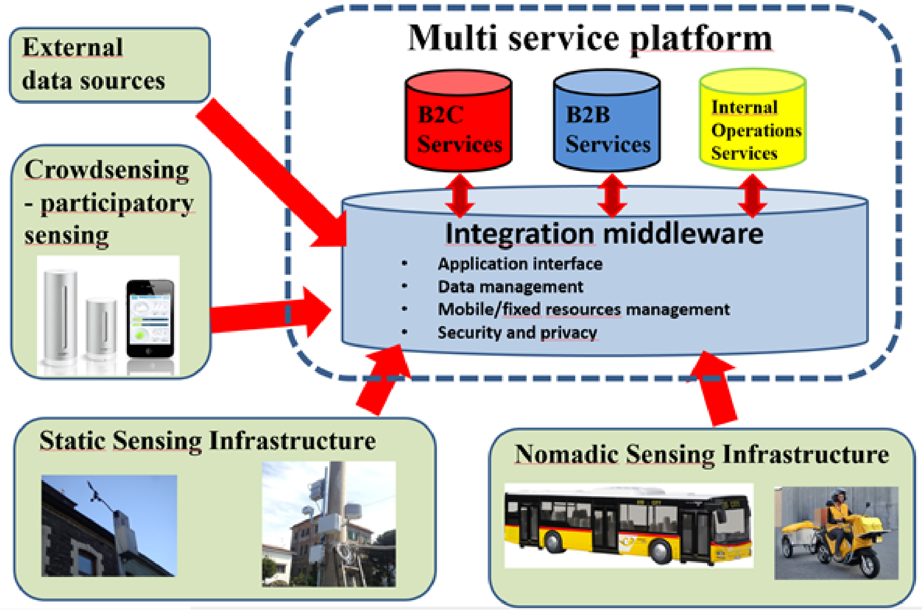

Fig. 1. Overview of the NOSE system architecture.

and on a measure of risk of accidents which combines data on road surface adherence with other metrics such as real time traffic information, past accident history, and so on. Such a service would potentially enable a higher efficiency of interventions for road maintenance, saving resources and time. It would enable targeted interventions of road maintenance, by helping to build a priority of interventions, based on their impact on safety.

3 The Architecture

3.1 Architecture design

In this section, we outline the architecture of our platform for nomadic sensing and computing, and its main functional blocks. The principles which inform the design of the architecture and of its components are:

- flexibility, and ease of (re-)configuration, in order to easily accommodate new services and data sources;
- resiliency, possibly through redundancy and autoconfiguration capabilities. For instance, it should be able to minimize the impact on application-level QoS of missing data due to temporary isolation of data sources from the Internet;
- simplicity, for cost effectiveness (and resiliency).

Fig. 3 shows a high level scheme of the architecture. A first type of functional blocks is represented by data sources, which can be broadly classified in:

- Nomadic. They consist in various environmental sensors, placed on vehicles of different type (scooters, buses, cars). We assume each vehicle to be endowed with GPS, so that we are always able to associate sensor readings to the spatial (and temporal) coordinates of the location where they have been collected.
- Static. They consist in data sources installed in fixed locations. Examples are weather stations and pollution sensors. Not being constrained by form factor, and energy considerations, they are typically more accurate, and enjoy a higher availability due, among others, to a more reliable connectivity.
- External data sources. They include external databases (e.g. open databases), as well as crowdsensed static and nomadic data, made available to the general public by private citizens, institutions and companies. They include temperature and wind sensors, webcams, etc.

At the core of the architecture there is an integration middleware, which incorporates several key functionalities, such as data aggregation and integration, security and privacy management, management of fixed and nomadic resources, QoS management, and a publish-subscribe system for event management. At the interface with applications, it exposes an API which facilitates development of services by hiding the complexity of the underlying system.

3.2 A first implementation

In this section, we describe the architecture of a first implementation of the NOSE platform, and the main implementation choices.

The overall architecture of the system is shown in Fig. 3.2. A first component of the system is the mobile data source. For the choice of the road surface monitoring sensor and environmental sensors, among the features we considered as desirable are support for outdoor deployment and harsh environmental conditions (low temperatures, up to -20° at least), support for measurements in movement, and small form factor for unobtrusive installation on small vehicles. Our resulting choice for road surface measurements has been the Marwis (Mobile Advanced Road Weather Information Sensor) sensor [5]. It is a mobile road weather information sensor that measures road conditions and environmental data. Unless otherwise specified, the sampling rate is $100Hz$. It collects the following data:

- Ambient temperature (range: -40°, 60°. Resolution: 0.1°);
- Dew point temperature (range: -50°, 60°, accuracy: 1.5°);
- Friction. Range: 0 (no friction) to 1 (perfect adherence);
- Relative humidity (range: 0% to 100%. Sampling rate: $10Hz$);
- Ice percentage (range: 0% to 100%);
- Road surface temperature. It is measured using a pyrometer. The measurement range is -40° to 70°, with a resolution of 0.1° and an accuracy of plus or minus 0.8° at 0°. Measurements are taken at a sampling rate of $10Hz$.
- Water film height (range: 0 to $6000\mu m$. Resolution: $0.1\mu m$).

Of all measurements, the most interesting from a driver's perspective is road friction. Note that it is not the result of a direct measurement, but it is an estimation made on the basis of a physical model, parametrized with all the other measurements made by the sensor. However, comparative studies have shown that the friction model used by the Marwis sensor is very accurate over a wide range of environmental conditions [8]. All the other measurements made by the Marwis

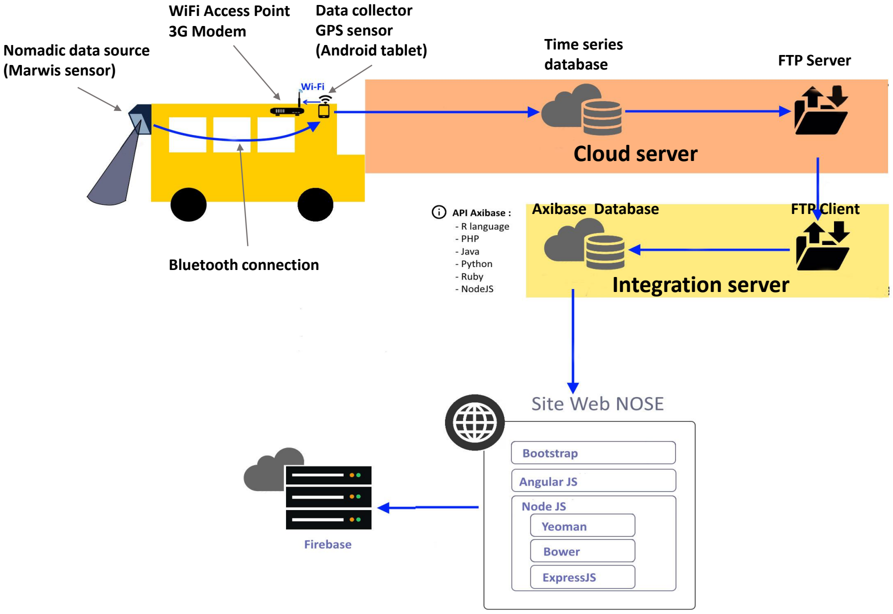

Fig. 2. Detailed architecture of a first implementation of the NOSE system.

are however useful to determine the factors which determine a given value of grip, and hence their correlation with a given level of risk of accident.

The Marwis sensor has been installed on a public transportation bus, and connected via Bluetooth to an Android tablet, residing in the same vehicle as the sensor. On such a tablet, an app takes care of collecting data from the sensor, of estimating parameters such as friction, for which there is no direct measurement available. The same app is also responsible for reading GPS positioning data, and of coupling them with sensor readings. Such space and time series are then sent to a cloud server via a 3G connection, for backup purposes. The Android app takes care of managing the connection with the time series database residing in the cloud, buffering data when 3G connectivity is not available.

From the cloud server, data is periodically transferred to an Axybase time series database server, whose role is to consolidate and integrate all data coming from all data sources in the system. Due to technical issues in the cloud server, the synchronization between the two databases has been implemented via an FTP server. A script has taken care of periodically retrieving data from the FTP server and transferring it to the integration server. This has generated a synchronization delay of a few minutes between the two time series databases.

From the Axybase server, data is transferred to a web server, using JSON and various RESTful calls. The web site is implemented using the AngularJS JavaScript framework [11], with Bootstrap [17] as CSS standard. Given that we assumed that the majority of the users of the services provided by the NOSE platform would access it from a smartphone, we have adopted Firebase [3] as a web application platform.

4 Assessment

It is clear from the description of the architecture that the principle of simplicity and cost effectiveness has informed all the main choices made. Indeed, in our implementation choices we have tried to maximize the reuse of existing solutions and building blocks, for the sake of minimizing integration and engineering costs, and maximizing reliability.

Nonetheless, system reliability has been one of the main challenges we have faced in the implementation of our system. The core issue we have faced has been represented by the fact that none of the components of the system we have used has been conceived for continuous, uninterrupted use, and for use in an environment, such as the one of a vehicle, characterized by mechanical stress and vibrations, and by instability in the supply of power and connectivity.

Thee most reliable part of the system has been the Marwis sensor, which has been built for continuous operation. The Bluetooth connection between the Marwis sensor and the Android app for data collection has experience periodic failures, despite the position of the two devices never changed over time during the whole evaluation. These failures have been probably due to the changing propagation conditions inside the bus, due to presence of passengers, to their position in the bus and their density. As expected, the system composed by a tablet/smartphone and the Marwis app has never run continuously for more than a few days, without requiring some form of intervention (such as reboot/app restart). The high variability in the operating hours of the bus within the week has entailed several challenges for the autonomous operation of the system. Hence, to save energy when the bus is not moving, we have implemented a few scripts which automatically turned off the network interface on the tablet, for the duration of the inactivity period. Finally, as expected in an mountain environment, 3G connectivity has been discontinuous and spotty, occasionally originating some data loss.

5 Conclusions

To date, the opportunities offered by moving infrastructure for realizing the potential of smart city paradigm, despite having been clearly acknowledged and made the object of several research projects (see section on the state of the art), remain largely unexploited. Among the possible explanations for this, we have the relative novelty of the smart city paradigm, which, already in its most classical forms, has started to take over only in these recent years. Moreover, far from being a technical solution, the smart city paradigm requires a new way of thinking about the urban setting and all the processes taking place in it, which is not yet there. Finally, several technical issues related to the management of a mobile, multi-service sensing, computing and communication infrastructure are still open.

In this paper we have presented an approach for low-cost nomadic smart city infrastructure, which aims at minimizing engineering and maintenance by the use of COTS solutions, and which decreases the amount of installed infrastructure by exploiting the ”infrastructure on wheels” existing in a district. Implementation and first assessment have shown the high impact of maintenance and configuration costs on the total cost of the solution, due to the difficulty in achieving high levels of infrastructure availability in such a volatile environment as onboard of public transportation buses.

Acknowledgements

This publication was produced in the framework of the RCSO NOSE project, under HES SO agreement 51668. The authors wish to thank PostAuto- CarPostal for its support.

References

1. European academic smart cities network (2014), `http://www.corp.at/archive/ppt/CORP2014_SLIDES174.pdf`
2. Federal statistical office, switzerland (2017), `https://www.bfs.admin.ch`
3. Firebase mobile app platform (2017), `https://firebase.google.com`
4. First mile solutions (2017), `http://www.firstmilesolutions.com`
5. Marwis specifications (2017), `{http://www.lufft-marwis.de}`
6. Mit senseable city lab (2017), `senseable.mit.edu`
7. Smart santander (2017), `www.smartsantander.eu/`
8. Bunnell, W., Krohn, D., Mathew, J.K., McNamara, M., Mekker, M., Richardson, L., Bullock, D.M.: Evaluation of mobile advanced road weather information sensor (marwis) by lufft for indiana winter road operations (2016), `http://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=4025&context=roadschool`
9. Deakin, M.: The intelcities community of practice: the egov services model for socially inclusive and participatory urban regeneration programs. In: Web-Based Learning Solutions for Communities of Practice: Developing Virtual Environments for Social and Pedagogical Advancement, pp. 263–284. IGI Global (2010)
10. Eriksson, J., Girod, L., Hull, B., Newton, R., Madden, S., Balakrishnan, H.: The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring. In: The Sixth Annual International conference on Mobile Systems, Applications and Services (MobiSys 2008). Breckenridge, U.S.A. (June 2008)
11. Green, B., Seshadri, S.: AngularJS. ” O’Reilly Media, Inc.” (2013)
12. Hull, B., Bychkovsky, V., Zhang, Y., Chen, K., Goraczko, M., Miu, A., Shih, E., Balakrishnan, H., Madden, S.: Cartel: a distributed mobile sensor computing system. In: Proceedings of the 4th international conference on Embedded networked sensor systems. pp. 125–138. ACM (2006)
13. Mertz, C.: Continuous road damage detection using regular service vehicles. In: Proceedings of the ITS World Congress (2011)

14. Paradells, J., Gomez, C., Demirkol, I., Oller, J., Catalan, M.: Infrastructureless smart cities. use cases and performance. In: Smart Communications in Network Technologies (SaCoNeT), 2014 International Conference on. pp. 1–6. IEEE (2014)
15. Schaffers, H., Komninos, N., Pallot, M., Aguas, M., Almirall, E., Bakici, T., Barroca, J., Carter, D., Corriou, M., Fernadez, J., et al.: Smart cities as innovation ecosystems sustained by the future internet (2012)
16. Soroush, H., Banerjee, N., Balasubramanian, A., Corner, M.D., Levine, B.N., Lynn, B.: Dome: a diverse outdoor mobile testbed. In: Proceedings of the 1st ACM International Workshop on Hot Topics of Planet-Scale Mobility Measurements. p. 2. ACM (2009)
17. Spurlock, J.: Bootstrap. " O'Reilly Media, Inc." (2013)