

# CityPulse: Real-Time IoT Stream Processing and Large-scale Data Analytics for Smart City Applications

Payam Barnaghi<sup>1</sup>, Ralf Tönjes<sup>2</sup>, Jan Höller<sup>3</sup>, Manfred Hauswirth<sup>4</sup>, Amit Sheth<sup>5</sup>, Pramod Anantharam<sup>5</sup>

<sup>1</sup> University of Surrey, United Kingdom

<sup>2</sup> University of Applied Sciences Osnabrück, Germany

<sup>3</sup> Ericsson Research, Sweden,

<sup>4</sup> National University of Ireland, Galway, Ireland

<sup>5</sup> Wright State University, USA

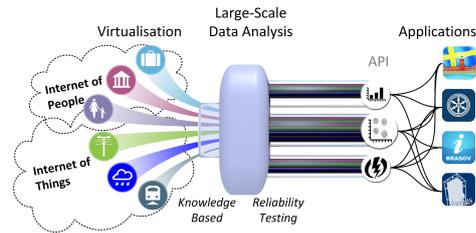
**Abstract.** The CityPulse project focuses on developing, building and testing a distributed framework for semantic discovery and processing of large-scale real-time IoT and relevant social data streams for knowledge extraction in a city environment. The project facilitates the creation and provision of reliable real-time smart city applications by bringing together the two disciplines of knowledge-based computing and reliability testing.

## 1 Introduction

The smart cities are evolving into a larger ecosystem or ecosystems that were previously disconnected. More and more applications and services in these ecosystems are going on-line. Today large amounts of valuable data and sensor information remain unused or are limited to specific application domains due to the large number of specific technologies and formats (such as traffic information, parking spaces, bus timetables, waiting times at events, event calendars, environment sensors for pollution or weather warnings, GIS databases). The aggregation of information from various sources is also typically done manually and the shared data is often static. The CityPulse project (<http://www.ict-citypulse.eu>) focuses on providing large-scale stream processing solutions to interlink data from Internet of Things and relevant social networks and to extract real-time information for the sustainable and smart city applications. The project supports the integration of dynamic data sources and context-dependent on-demand adaptations of processing chains during run-time. An overall view of the CityPulse framework is shown in Figure 1.

## 2 R&D activities and goals

Smart city data is big data. It is not only large in volume, it is multi-modal, varies in quality, format, representation form and levels of dynamicity. The data needs



**Fig. 1.** An overall view of the CityPulse framework

to be processed, aggregated and higher-level abstractions need to be created from the data to make it suitable for the event processing, knowledge extractions and event processing applications that enable intelligent applications and services for smart city platforms [1]. Data needs to be integrated from various domains and the resulting knowledge exposed to various domains in a federated fashion. The CityPulse framework is organised in three consecutive iteratively applied processing layers, covering federation of heterogeneous data streams, large-scale IoT stream processing, and real-time information processing and knowledge extraction.

### 3 Relevance to the session

CityPulse provides enrich data streams from physical and virtual sensing devices with semantic annotations, enabling adaptive processing, aggregation and federation of data and will provide actuation mechanisms to interact with the environment. Functionalists for aggregation and federation using linked data and mash-ups techniques assure a scalable framework for processing large-scale IoT data streams. One of the activities in the project is virtualisation of the Internet of Things resources and semantic annotation of data and streams. In this context, we are interested in using existing common ontologies such as W3C SSN [2] and other data annotation frameworks to describe the streams and their providing resources as well quality related features of the data. Participating in this session and discussing and sharing the ideas and experiences with the other projects will help us to avoid duplicate work, re-use and extend existing models and discuss possible collaboration with other projects.

### References

1. Barnaghi, P., Sheth, A., Henson, C., "From Data to Actionable Knowledge: Big Data Challenges in the Web of Things", *Intelligent Systems, IEEE* 28(6), pp. 6-11, 2013.
2. Compton, M, et al., "The SSN Ontology of the W3C Semantic Sensor Network Incubator Group", *Journal of Web Semantics*, 17. pp. 25-32, 2012.