

GRANT AGREEMENT No 609035
FP7-SMARTCITIES-2013

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



Collaborative Project

Final Report

Document Ref.	D1.4
Document Type	Report
Workpackage	WP1
Lead Contractor	UNIS
Author(s)	CityPulse Consortium
Contributing Partners	AA, AI, ERIC, NUIG, SAGO, SIE, UASO, WSU, UniS
Planned Delivery Date	M37
Actual Delivery Date	M38
Dissemination Level	Public
Status	Final
Version	v1.0
Reviewed by	Payam Barnaghi, Ralf Toenjes

Executive Summary

Cities are facing many economic, ecological and social challenges. To address these challenges better information and closer cooperation between the involved stakeholders is the key. This means smart technology and smart organisation. Technology is getting cheaper and smaller and is embedded into a city's infrastructure and citizens' smartphones connecting an Internet of things and people. However, the heterogeneity of the different data sources and missing integration platform hinder the uptake of innovative cross-domain smart city applications. The provided information (e.g. transportation timetables) is mainly static not reflecting the dynamics in a city or often just not found.

CityPulse tackled these issues in order to help municipalities and developers in creating better city services. The project implemented, integrated and tested a flexible and extensible smart city data analytics framework. Working with the city partners in Aarhus and Brasov, the project team developed software components for real-time data stream publication and annotation, accesses and query interfaces, semantic reasoning and information processing mechanisms, adaptive analysis methods for pattern recognition and correlation analysis, and social media data analysis. The project's efforts have also included creating awareness, collecting the requirements, defining use-cases and building the technical infrastructure. The project team, in collaboration with citizens, stakeholder groups, and the technical team, defined over 101 smart city use-case scenarios that were published and made available online. Showcasing the results, working with our project partners and receiving feedback at the development stage allowed the project partners to better understand the priorities of the citizens and city authorities. Dependability and reliability of the proposed solutions were also key issues taken into account. The project team learned several valuable lessons by involving a multi-disciplinary team in our development efforts.

During the lifetime of the project, the CityPulse team developed 19 open-source tools and software, defined and published 101 smart city scenarios, contributed to and participated in over 50 dissemination events, trained 6 PhD students, and published 11 journal articles and 26 conference papers. The partners have also commercialised extended versions of two of the developed solutions and the project's industrial partners will use the results of the projects in their future demonstrators and products.

This report summarises the key activities, development and results of the project.

www.ict-citypulse.eu

Table of contents

1	Project Objectives and Key Issues	6
1.1	Introduction.....	6
1.2	Project Objectives and Approach.....	7
1.3	Key Topics and Achievements.....	11
2	Main Scientific and Technical Results	15
2.1	Overview	15
2.2	Scenarios and Smart City Framework	19
2.2.1	Scenarios and Requirements.....	19
2.2.2	Smart City Framework	20
2.2.3	Reference Datasets for Performance Evaluation.....	22
2.3	Large-Scale IoT Stream Processing	24
2.3.1	Resource Virtualisation and Semantic Annotation.....	24
2.3.2	Data Federation and Aggregation.....	24
2.3.3	Event Detection for Urban Data Streams.....	25
2.3.4	Knowledge Acquisition Toolkit (KAT).....	26
2.3.5	Resource Management.....	27
2.4	Reliable Information Processing	28
2.4.1	Accuracy, Trustworthiness and QoI	28
2.4.2	Conflict Resolution and Fault Recovery.....	29
2.4.3	Monitoring and Testing.....	30
2.4.3.1	Atomic Monitoring.....	30
2.4.3.2	Composite Monitoring	33
2.4.3.3	Testing.....	37
2.4.4	QoI Explorer Open Source Tool.....	38
2.4.5	Event Test App.....	38
2.5	Real-Time IoT Intelligence	39
2.5.1	Real-Time Adaptive Urban Reasoning	40
2.5.1.1	Technical Adaptation in Stream Processing	40
2.5.1.2	Contextual Filtering	41
2.5.2	User-centric Decision Support in Dynamic Environments	42
2.5.3	Visualization in Dynamic Smart City Environments	43
2.5.4	Decision Support and Contextual Filtering Open Source Tool	44
2.6	Integration & Evaluation: Smart City Applications.....	45
2.6.1	Evaluation Scenarios	45
2.6.2	Overall Evaluation	46

3	Project Impact, Dissemination and Exploitation of Results.....	51
3.1	CityPulse Value Proposition in the Context of Smart Cities.....	51
3.2	Impact.....	53
3.3	Open Source Tools	55
3.4	Standardisation	59
3.5	Dissemination of Results.....	60
3.6	Public website	60
3.7	Beyond CityPulse	61
3.8	The CityPulse Consortium	62
4	Use and Dissemination of Foreground	63
4.1	Publications	63
4.1.1	Journal.....	63
4.1.2	Conference.....	64
4.2	Dissemination Activities	66
4.2.1	Exploitable Foreground	74
5	Conclusions.....	78
	References	79

List of Figures

Figure 1: Integrated Approach of CityPulse	7
Figure 2: Processing Steps during Different Life-cycle Stages.....	9
Figure 3: CityPulse Overview.....	11
Figure 4: CityPulse Demonstrators	16
Figure 5: CityPulse 3D Map Demo	16
Figure 6: CityPulse Social Media Analyser	17
Figure 7: CityPulse Quality Explorer	17
Figure 8: CityPulse Journey Planner Mobile App	18
Figure 9: CityPulse Smart City Dashboard	18
Figure 10: Top 10 Ranked Use-case Scenario Evaluations.....	20
Figure 11: CityPulse Architecture and APIs [Puiu, et. al, 2016].....	21
Figure 12: CityPulse Datasets.....	23
Figure 13: Virtualisation and Real-time Annotation in Smart City Applications	24
Figure 14: Architecture of the Data Federation and Aggregation Component	25
Figure 15: Event Detection Node.....	26
Figure 16: KAT Data Analytics Toolkit	27
Figure 17: Reliable Information Processing Overview.....	28
Figure 18: Quality Ontology	29
Figure 19: Fault Recovery Component Workflow	30
Figure 20: Atomic Monitoring Functionality.....	31
Figure 21: Determination of the QoI Metric Frequency.....	31
Figure 22: Update intervals for the traffic data in Aarhus	32
Figure 23: Aarhus Data Frequency Map.....	32
Figure 24: Activity Diagram of the Composite Monitoring Process.....	33
Figure 25: Example of Space-Time Resolution of Sensors and Events.....	34
Figure 26: Avg. Speed (Irregular) and Vehicle Count (Raw) Change During Traffic Incidents	35
Figure 27: Heat Map of Measured Value Pairs of Measured Traffic Speed Changes.....	36
Figure 28: Execution of Test Cases.....	37
Figure 29: Event Report Mobile Test App.....	38
Figure 30: Overview of the Real-time IoT Intelligence Components	39
Figure 31: Contextual Filtering Interactions	41
Figure 32: Event-based User-Centric Decision Support.....	42
Figure 33: Decision Support Interactions.....	42
Figure 34: CityPulse GitHub.....	59

1 Project Objectives and Key Issues

The CityPulse project started in September 2013 as a 3 years' project with 10 consortium members. The project's main goal was to provide a large-scale data analytics framework for smart cities to accelerate the introduction of smart city applications. The project's key goals were defined as:

- Virtualisation and hiding the heterogeneity of the numerous data and information sources,
- Large-scale data analytics for resource efficient event detection in multiple data streams,
- Semantic description frameworks and semantic analytics tools to provide machine-interpretable descriptions of data and knowledge across domains and at different abstraction levels,
- Easy creation of real-time smart city applications by re-usable intelligent components.

This report summarises the key efforts and provides an overview of the outcome and results of the project.

1.1 Introduction

An increasing number of cities have started to introduce new Information and Communication Technology (ICT) enabled services with the objective of addressing sustainability as well as improving the operational efficiency of services and infrastructure. In addition, there is increased interest in providing novel or enhanced service offerings and improved experiences to citizens and businesses. These cities are thus evolving into a larger eco-system or ecosystems that were previously more or less unconnected. More and more of the activities in these ecosystems are going on-line. The city councils are the pivotal facilitators in making the on-line ecosystem of ecosystems become a reality.

However, a challenge in the smart city approach is integration across different application domains, as well as the engagement of different city departments, city-contracted entrepreneurs and individual enterprises providing services. 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 (traffic information, parking spaces, bus timetables, waiting times at events, event calendars, environment sensors for pollution or weather warnings, GIS databases etc.). Hence, an aggregation of information from various sources is typically done manually and is often out-dated or just static. The smart cities vision is hindered by the following problems:

- The heterogeneity of the various data sources and the missing integration platform hinder the uptake of innovative cross domain smart city applications.
- The large amount of raw data without intrinsic explanation remains meaningless in the context of other application domains, thus hindering automated integration in innovative applications.
- The information is static and does not reflect the actual state of systems and dynamics in the city (e.g. people moving within a city in most cases base their decisions on printed transportation timetables, static opening hours, etc.). Mechanisms that provide real-time data to smart city applications and citizens are missing.
- The data generated is large both in the sheer number of sources as well as in time, i.e. The sources include a large number of data streams. However, each data item is very small. What is missing are scalable tools to provide aggregation and knowledge extraction from these large sets of data streams and match them with the needs of the applications.
- Both data and extracted knowledge need to be imported and exported, thus requiring tools for federation. The appropriate means of providing meaning to different abstraction levels of Internet of Things (IoT) data and knowledge do not exist.
- The most valuable resources of a smart city are its infrastructure and its citizens. However, the IoT and the Internet of People (IoP), i.e. social media, are not integrated information wise. This requires methods for integrating IoT and social network data that pay particular attention to the accuracy and trustworthiness of the data.

- Data fusion and information integration from sensory and citizen data sources are often planned at design time and current systems lack real-time, online and dynamic integration, processing, adaptation and visualisation of the various smart city data streams.

Bridging technology and domain boundaries, thus enabling a wider uptake of smart city applications, requires a scalable, adaptive and robust platform that supports the integration and analysis of heterogeneous data and information sources and facilitates the development of innovative real-time smart city applications.

The Internet of Things is an error prone environment. Therefore, CityPulse has integrated reliability testing capabilities in IoT application development and provision from the beginning. This is particularly important when IoT is combined with Internet of People, where trust is a major issue. The project has investigated how formal test procedures can be integrated into the development of the application and especially into the creation environment of the application by using a systematic approach.

The achievement of sustainability for a smart city platform is challenged by the dynamics of the evolving city ecosystems. This project has addressed this issue by using a knowledge-based approach to facilitate the maintenance, expandability and use of platform components. Semantic models and semantic analytics support a goal-driven configuration and adaptation of the CityPulse framework components. CityPulse has enabled innovative smart city applications by adopting an integrated approach to the Internet of Things and the Internet of People and by integrating Physical-Cyber-Social data and providing analytical tools and mechanisms for processing multi-modal and dynamic real world data (see Figure 1).




Figure 1: Integrated Approach of CityPulse

1.2 Project Objectives and Approach

The main objective of the project was to:

Develop, build and test a distributed framework for the semantic discovery and processing of large-scale real-time IoT and relevant social data streams for knowledge extraction in a city environment.

To achieve this objective, the project has developed a large set of software components, architecture and best practices, use-case-scenarios and demonstrators that are presented as the CityPulse framework in this report. The software components and documentations are made publically available and source codes are provided as open-source under the MIT License. The CityPulse project has demonstrated the integration of dynamic data sources and context-dependent on-demand adaptations

of processing chains during run-time. The developed tools and components aim to bridge the gap between the application technologies on the IoT and real world data streams. They address the specific requirements of business services and user applications in smart city environments that use Cyber-Physical and Social data and employ big data analytics and intelligent methods to aggregate, interpret and extract meaningful knowledge and perceptions from large sets of heterogeneous data streams.

The framework includes middleware, common interfaces and semantic models, and different components and processing methods that enable smart city applications using human and machine sensory data. The components and interfaces are made open and flexible to support development of third-party applications and services using the CityPulse framework. The project has also developed a large number of smart city use-cases in collaboration with city partners and city stakeholder group members. The use-cases focus on the key issues and describe narratives and requirements for each particular scenario. The information is made available online on the project website (<http://ict-citypulse.eu>) and described in related project reports.

Smart city data is big data. It is multi-modal and varies in quality and format and representation form. The data needs 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. Data needs to be integrated from various domains and the resulting knowledge exposed to various domains in a federated fashion.

CityPulse has provided large-scale stream processing solutions to interlink data from Internet of Things and relevant Social Networks (i.e. Twitter and social media analysis) to extract real-time information for the sustainable and smart city applications. The sources of data streams considered in the CityPulse framework are:

- Fixed sensors in the city infrastructure (e.g. streets, public buildings, utility systems) or in personal and business properties (e.g. vehicles, homes, buildings)
- Citizen owned mobile phones captured through opportunistic and participatory sensing
- Social streams/feeds filtered and collected by agents from Facebook, Twitter, blogs or other social media platforms
- City information systems providing high-level information (e.g. GIS)

The project team has developed various solutions that are briefly described in this report and has published the results of their evolution and has made the data available for other researchers and developers in the community to develop new solutions and compare their results with the CityPulse outcomes. The project team hopes this will encourage development of enhanced applications and analytics solutions for smart city applications and will also provide a baseline for comparing and evaluating the research and development results. The set of key datasets are published as CityPulse Reference Datasets that consists of over 180GB of time series data and also live data feeds. The reference datasets are available via CityPulse website at: <http://iot.ee.surrey.ac.uk:8080>

The data analytics, semantic reasoning and machine learning research in the project mainly focused on real-time and adaptable data analytics, reasoning and information extraction and quality analysis for city data streams to provide knowledge-enabled reliable and sustainable smart city applications while considering the privacy and integrity of individuals and businesses. 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. To achieve reliability, CityPulse integrates data analytics and reasoning

methods with reliability monitoring and testing at all stages of the data stream processing and interpretation. The latter provides solutions for the different life-cycle stages of data processing and utilisation, supporting application development, i.e. design-time, and application provision, i.e. run-time. The processing steps and related objectives are shown in Figure 2 and described as follows:

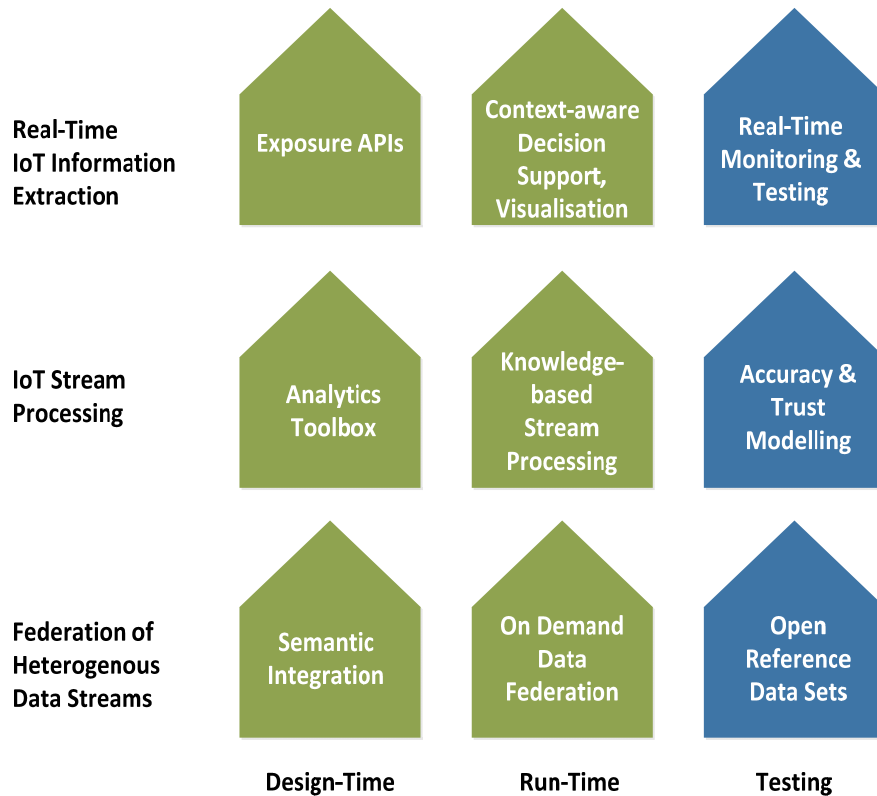


Figure 2: Processing Steps during Different Life-cycle Stages

Description of the processing steps related to the life cycle:

- Federation of heterogeneous data streams
 - Design-time:
 - Mechanisms for easy integration and publication of new data streams from IoT devices, social media, and information systems.
 - Syntax and semantics for annotation of data streams enabling automated search, discovery and federation of data streams.
 - Support for open data trading.
 - Run-time:
 - Semi-automatically and semantically driven data stream discovery and federation from heterogeneous sources.
 - Adaptive integration using linked data and mash-ups of Cyber-Physical-Social data streams (i.e. IoT and relevant social media streams).
 - Reliability testing:
 - Reference data sets for performance evaluation that are made available to the IoT community.
- Large-scale IoT stream processing
 - Design-time:
 - Toolbox with re-usable components for big data analytics (including data aggregation and summarisation, pattern matching, data mining, distributed stream processing, mash-ups, event detection, and pattern learning).

- Sustainable middleware allowing evolutionary extension of toolbox.
- Run-time:
 - Knowledge-based methods for automated aggregation and abstraction of semantically annotated IoT data streams using big data analytics tools.
 - Goal-driven configuration and dynamic adaptation of IoT stream processing components and parameters during run-time.
- Reliability testing:
 - Methods for modelling and processing uncertain and imprecise data.
 - Methods for modelling and processing trust for linked Cyber-Physical-Social data.
- Real-time IoT information processing and knowledge extraction
 - Design-time:
 - Modelling abstractions and Application Programming Interface (API) for smart city application developers.
 - Toolbox with re-usable components for real-time IoT information extraction (including knowledge-based interpretation, complex urban reasoning, matchmaking, context dependent information extraction, and decision support).
 - Methods and tools for knowledge modelling enabling explicit control of component adaptation in dynamic smart city environments
 - Run-time:
 - Monitoring and knowledge based control and adaptation of IoT data and information processing components in a dynamic IoT environment.
 - Context dependent and situation-aware user-centric extraction of real-time IoT information.
 - Visualisation of real-time urban IoT information.
 - Reliability testing:
 - Tools for real-time monitoring of the information extraction process
 - Definition and implementation of benchmark test to compare and evaluate the alternative solutions for real-time smart city applications.

CityPulse combines the above sources of information and intelligent techniques to offer suitable modelling abstractions and APIs for:

- Data analysts/knowledge engineers with domain expertise to provide knowledge/event plug-ins that can generate new insights from Cyber-Physical-Social data streams;
- Application developers and service providers that allow the creation of innovative smart city services based on the combination of information flows coming from knowledge/ event plugins or raw IoT data streams;
- Citizens and businesses to search and discover data related to public infrastructure, events and experiences, and subscribe to update/event notifications related to events.

The project has developed proof-of-concept and has prototyped its major concepts. The project has also defined a comprehensive set of metrics and has provided an extensive set of evaluations of the results of the experiments and developments. The framework components are also validated in different smart city application scenarios, such as intelligent combined public and individual transport, large event analysis, and resource monitoring and planning.

To achieve these challenging objectives the project brought together a well-focused consortium of city authorities, industry and academia with expertise in Internet of Things, semantic data processing and testing. The partners built upon an extensive track record of developing IoT solutions in previous

projects and designed the key architecture and system models compatible and based on common architectures and models in the IoT domain. The collaboration with multiple city authorities (Aarhus and Brasov as project partners, Stockholm and Osnabrück as non-founded collaborative partners) provided access to extensive ideas and expertise in the smart city and technical and non-technical problems facing the cities. This led to design and development of real-world problem focused solutions and helped the project team to define a large-set of scenarios and use-cases and specify the key requirements and analyse them. This pragmatic approach and citizen and city expert guided approach have significantly contributed to the success and popularity of the project results including the requirements analysis, use-case, toolkits and open-source components (please visit the project website and the reports available at: <http://www.ict-citypulse.eu>).

1.3 Key Topics and Achievements

CityPulse provides novel approaches to analyse dynamic IoT data streams and social data streams, by using intelligent methods for big IoT data analytics and smart city services and processes. Figure 3 depicts an architectural overview of the CityPulse topics and key components that address the project's key issues.

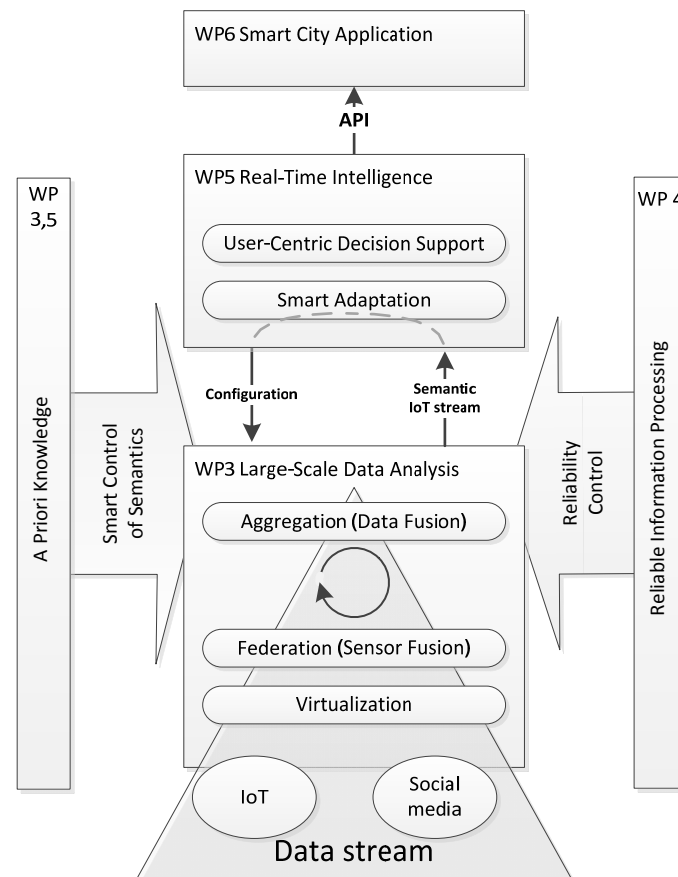


Figure 3: CityPulse Overview

The Key issues addressed in CityPulse (described in relation to the components shown in Figure 3) are described below.

Virtualisation: Semantic annotation of heterogeneous data for automated discovery and knowledge-based processing

A major hindrance of the uptake of Smart City applications is the heterogeneity of the data and underlying networks and resources. As a result, the applications are usually data specific, often

assuming implicit knowledge about the data and limiting the possibility of reusing the data for other applications. To overcome these silo architectures CityPulse developed an integrated framework. CityPulse systematically applies virtualisation of the underlying resources. The virtualisation provides a common abstract interface to the higher layers that deal with information processing and knowledge extraction. In this project data virtualisation is offered to data consumers as a common access interface that hides the technical aspects of data streams, such as location, storage structure, access language, and streaming technology.

The project has also developed a set of semantic annotation model and stream annotation methods and analytical solutions that provide interoperability in heterogeneous environments. The models include Stream Annotation Ontology (<http://purl.oclc.org/NET/UNIS/sao/sao#>), Quality Ontology (<http://purl.oclc.org/NET/UASO/qoi>), and stream annotation tools and methods including SAOPY libraries (<http://iot.ee.surrey.ac.uk/citypulse/ontologies/sao/saopy.html>) and Ontology Validation service (<http://iot.ee.surrey.ac.uk/SSNValidation/>).

This layer is the key enabler to ensure easy discovery and automated processing of data. CityPulse approaches this problem by semantic annotation of the data streams, thus making their characteristics and capabilities available in a machine process-able way. The meta-data for the data streams is employed to automate the discovery and process control for data analytic [Kolozali et al., 2016]. The project has also developed distributed mechanisms to index and search large-scale annotated data [Fathy et al., 2016], and to process time-series data [Ganz et al., 2015], both for numerical [Puschmann et al., 2016], [Puschmann et al., 2016-2] and also social data-streams [Anantharam et al., 2015], [FarajiDavar et al., 2016].

Federation: On demand integration of heterogeneous Cyber-Physical-Social sources

Data federation is a form of data virtualisation in which heterogeneous data streams are made accessible to data consumers as one integrated data stream by using on-demand data integration. This component combines heterogeneous sets of data streams to create one unified view of all the data. The IoT environment and the physical world are subject to constant changes and the devices that provide the underlying layer for IoT are often resource constrained, pervasive, unreliable and heterogeneous in nature. This dynamicity and complexity hinders the effective utilisation of existing data mining and information processing technologies in the IoT world and makes it challenging to effectively and seamlessly integrate the real world data with the existing information systems and data analytics solutions. Current integration of data streams is typically done on the basis of static artefacts (e.g. documentation of available data streams and services) and changes in underlying resources require manual changes in existing federations.

CityPulse has developed adaptive and flexible methods taking into the account the dynamic data coming from the urban infrastructure, human sensors, social web data and other data sources [Puiu et al., 2016], [FarajiDavar et al., 2016], [Puschmann et al., 2016]. CityPulse also provides a knowledge-based approach using semantics and linked data to automate the data federation [Kolozali et al., 2014]. A linked-data approach is employed to represent data stream relationships and to support mash-ups [Gao et al., 2016], [Gao et al., 2015]. The focus has been on developing adaptive methods that can cope with real-time federation of multi-modal data streams [Puschmann et al., 2016], [Kolozali et al., 2016].

Aggregation: Large-scale data analytics

While the above federation provides a kind of sensor fusion, the aggregation aims at data fusion, thus reducing the enormous amount of data by aggregation techniques, e.g. clustering, summarisation, filtering and pattern recognition. The data is time and location dependent, is often transient and often needs to be processed in correlation with other data to create higher-level abstractions. Hence

aggregation and federation are applied consecutively and iteratively to extract the relevant data and to detect events. The aggregation and event detection provide machine-interpretable observations, i.e. semantically annotated streams, to the next higher layer interpreting the real world context. The challenge is the large scale of dynamic data.

CityPulse has developed develop solutions for large-scale data analytics including data aggregation and summarisation, pattern matching, data mining, distributed stream processing, mash-ups, event detection, and pattern learning. Part of the analytical methods and algorithms are integrated in a standalone and open-source toolkit, called KAT (<http://kat.ee.surrey.ac.uk>). KAT includes re-usable methods and algorithms that can be used in smart city data analytics [Ganz et al., 2015].

Smart Adaptation: Real-Time interpretation and data analytics control

CityPulse components are designed for dynamic smart city environments in which the properties of underlying resources and streams need to be constantly updated according to changes and events in the real world. In most of the existing solutions matchmaking between the requirements expressed by businesses and applications and available data is carried out at design-time. This approach is often far from optimal and its deficiencies become even more obvious in IoT scenarios where the properties of underlying services and resources dynamically change and depend on physical world events and phenomena (e.g. sensor readings - network availability, weather conditions, and temperature).

CityPulse provides solutions for adaptive data processing and utilisation and will enable knowledge-based automated composition of data/streams and resources using linked Cyber-Physical-Social resources [Barnaghi, et al., 2016], [Puiu et al., 2016]. By transforming the lower-level dynamic information (e.g., changes in sensor readings) to higher-level contextual abstractions, and using state-of-the-art semantic Web technologies (e.g., resource/knowledge descriptions frameworks, reasoning and rules), service matchmaking and business management processes can be realised in more efficient and meaningful ways. This layer provides methods for higher-level information processing to interpret the semantic data in the current context [Kolozali et al., 2014]. The data interpretation is challenged by the real-time demands of the targeted smart city applications. This requires the development of efficient methodologies that control what is processed in the fast big data analytics pipe and the knowledge available at the higher-level for interpretation. CityPulse employs domain knowledge to interpret the condensed data streams and detect higher-level events (i.e. machine interpretable or human understandable events) from Cyber-Physical-Social streams [Tönjes et al., 2015], [FarajiDavar et al., 2016].

User centric decision support: Context aware customized IoT information extraction

The goal of the user-centric decision support is to utilise contextual information, usage patterns and preferences to provide optimal configurations of smart city applications. The users from this perspective will be citizens, enterprises or city councils. User requirements and domain knowledge will be modelled using linked-data and open vocabularies in order to provide lightweight, interoperable and well-established foundations for decision making support and matchmaking of city services. Users will be able to specify their requirements and preferences explicitly and their requirements and preferences will be implicitly derived from their individual application usage patterns. The social and context analysis enables CityPulse matchmaking and discovery mechanisms to match the data according to the user's preferences and context-dependent attributes (location, time, etc.) [Fathy et al., 2016], [Anantharam et al., 2015-2], [Gao et al., 2015], [Nechifor et al., 2014].

Reliable Information Processing: Testing and monitoring accuracy and trust

The IoT systems often operate in dynamic environments that are subject to change and prone to errors. Developing reliable data processing and information extraction methods requires accuracy and trust

issues to be taken into consideration when dealing with IoT data. This is further emphasised when data from citizens (e.g. smart phone sensors or social media) is used.

CityPulse has developed measures and methods to include and process accuracy and trust in data acquisition, federation and aggregation. In this context an open-source toolkit called Quality Explorer is developed (<http://131.227.92.55:8014/>). The Quality Explorer provides machine readable and semantically annotated quality analysis of the online data streams. The quality explorer also provides a visual interface for end-users (e.g. city planners) to browse and see the quality related information for the underlying sensory data collection infrastructure on smart city environments. Various quality assessment, analysis and compensation mechanisms are implemented in this component. This provides a set of integrate techniques for monitoring (run-time) and testing (design-time) in the framework, thereby ensuring reliable information processing.

Smart City Applications: Application programming interface for rapid prototyping

Many of today's cities are faced with challenges resulting from the changes and new demands that a rapidly growing digital economy imposes on current applications and information systems. In order to monitor, manage and plan for public resources and infrastructures in city environments, and to enable citizens and businesses to prosper and enjoy trustworthy and reliable data integration, efficient processing and analytics design, development and deployment are required; in the IoT domain in particular flexible solutions for sensing, extracting knowledge, processing large volumes of raw sensory data and responding to events in the physical world in real-time are required.

CityPulse provides a set of open-source components and demonstrators for to smart city application programmers, providing access to and management of the re-usable components described in the above sections. The components are available at the project main code repository (<https://github.com/CityPulse>).

2 Main Scientific and Technical Results

This section provides a summary of the key scientific and research results of the project and describes the main contributions to research and innovation in related domains.

2.1 Overview

The project has developed a set of open-source components, adaptable machine learning and data analytics algorithms for dynamic real world data and online time-series and social data analysis for smart city applications. The key achievements include:

Defining and presenting a large set of smart city use-case scenarios that are defined based on citizen requirements and city stakeholder group recommendations (<http://www.ict-citypulse.eu/scenarios>).

- A detailed architecture and requirement analysis for smart city-use cases (http://www.ict-citypulse.eu/page/sites/default/files/citypulse_d2.1_requirements_v1.0_0.pdf).
- A smart city framework architecture that was developed based on the IoT Architecture Reference Model (IoT-ARM) proposed by the European Research Cluster on the Internet of Things and the IOT-A project.
(http://www.ict-citypulse.eu/page/sites/default/files/citypulse_d2.2_smart_city_framework_v2.3.pdf).
- A large-set of reference data sets for smart city data analytics research and innovation evaluations (<http://iot.ee.surrey.ac.uk:8080/>).
- A set of lightweight and well-defined semantic and meta-data models, API and tools for virtualisation and description of large-scale IoT and social media data streams. (http://www.ict-citypulse.eu/page/sites/default/files/citypulse_d3.1_v1.3.pdf)
- Novel and adaptable machine learning and data analytics algorithms for analysing and processing dynamic smart city data (<http://kat.ee.surrey.ac.uk>).
- Algorithms for analysing and evaluating the quality of information for large-scale smart data and mechanisms for fault recovery (<https://github.com/CityPulse/Fault-Recovery>).
- A GitHub repository for the open-source software components and documentation for the CityPulse Smart City Framework. The codes and software are released under the MIT open-source license (<https://github.com/CityPulse>).
- Real-time analysis and information extraction algorithms and methods for social media data based on probabilistic machine learning and deep learning methods (<https://github.com/CityPulse/Social-Media-Analyser>).
- Novel solutions for semantic processing and large-scale stream reasoning for smart city applications (http://www.ict-citypulse.eu/page/sites/default/files/d5.1-citypulse_v1.8-final.pdf).
- Methods and apparatus for user-centric decision support in dynamic smart city environments (http://www.ict-citypulse.eu/page/sites/default/files/citypulse_d5.2_final.pdf).
- A set of demonstrators and open-source applications that are built based on the CityPulse framework components. Figure 4 shows a snapshot of some of the CityPulse demos and the following describes some of the developed applications.

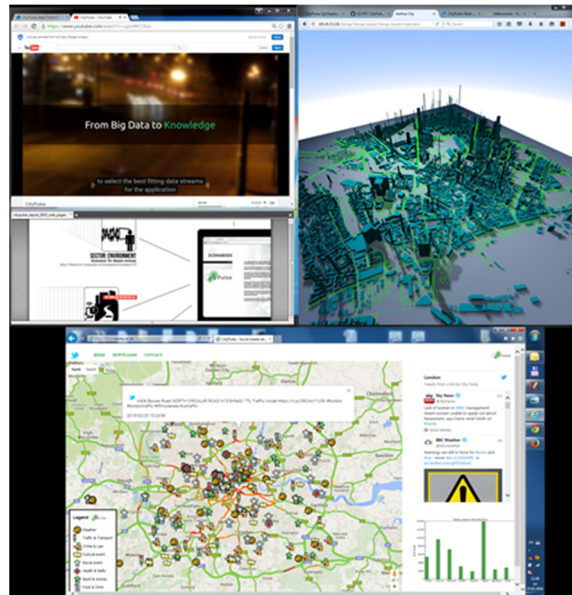


Figure 4: CityPulse Demonstrators

The project has identified 101 smart city scenarios and related use cases in cooperation with partner cities and city cooperation (City Stakeholder Group) and derived a set of requirements for a smart city framework based on proposed use cases, references in the field and “on site” workshops together with city partners. The requirements are encoded in an evaluation metric that has been made available online for the wider community to rank the 101 scenarios. The project team then developed a set of applications and demonstrators based on these scenarios and their requirements. Some of these demonstrators include:

- **3D Map application:** provides a 3D overview of the city of Aarhus (and a few other cities) and can visualise energy consumption in buildings. Figure 5 shows a screenshot of the 3D map application. There are two separate components: the WSServer and the 3Dmap - the first one is a web socket server that randomises buildings and energy consumptions (to preserve the privacy and security for public demo) and sends the information to all the connected clients. Another reason for having this component is due to the fact that the energy consumption of buildings in the city of Aarhus is not yet openly available in the Open Data portal of the city (<http://185.45.23.126/3dmap/3dmap-master/3dmap-master/index.html>).




Figure 5: CityPulse 3D Map Demo

- The social media (i.e. Internet of People) data processing component includes a dedicated Data wrapper unit, which uses the Twitter stream API, and Google translate API to simultaneously collect the tweets and to automatically detect the source language and translate the tweets to English. Figure 6 shows a screenshot of the demo. The data processing unit uses a Conditional Random Field Name Entity Recognition, a deep learning Convolutional Neural Network (CNN) for Part of Speech tagging and a multi-view event extraction to assign the twitter data into different smart city event categories. An online demo of the social media data analysis software is available at: <http://iot.ee.surrey.ac.uk/citypulse-social/>

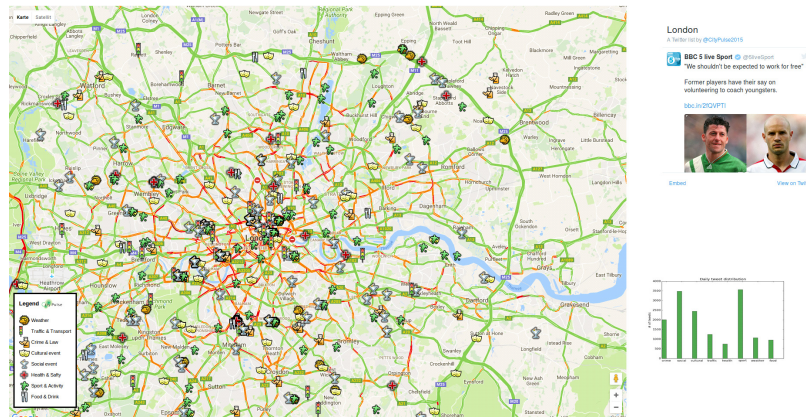


Figure 6: CityPulse Social Media Analyser

- Quality explorer: provides machine readable and semantically annotated quality analysis of the online data streams. The quality explorer also provides a visual interface for end-users (e.g. city planners) to browse and see the quality related information for the underlying sensory data collection infrastructure on smart city environments. Various quality assessment, analysis and compensation mechanisms are implemented in this component. Figure 7 illustrates the Quality Explorer applications and an online demo is available at: <http://131.227.92.55:8014/>

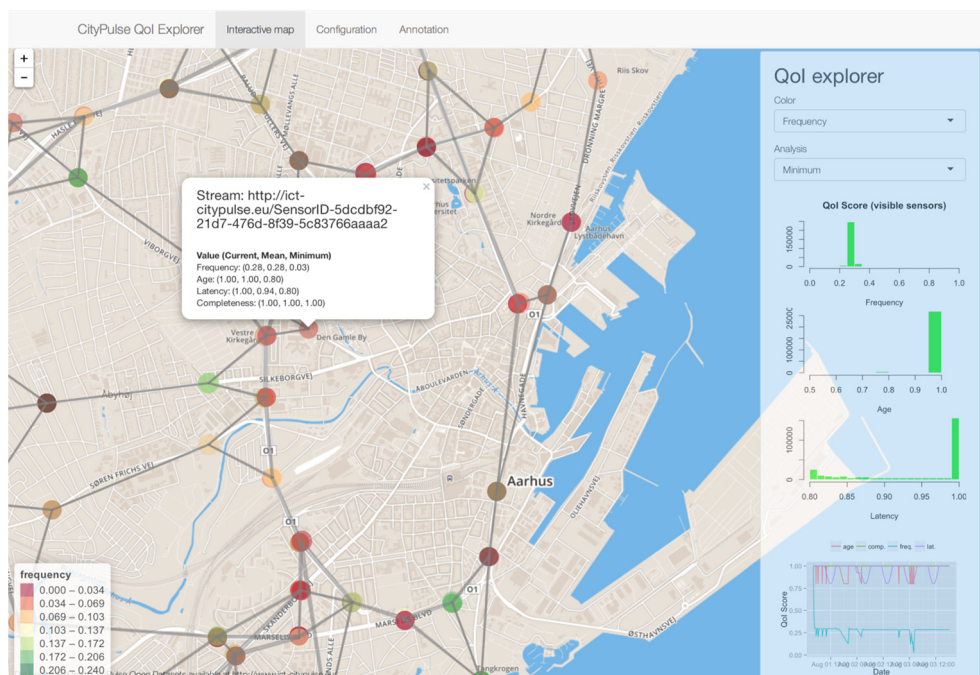


Figure 7: CityPulse Quality Explorer

- Journey planner mobile app (Figure 8) and smart city dashboard (Figure 9) are applications that are built on the existing components of the CityPulse framework and demonstrate how an integrated set of CityPulse data analytics components can be used to create smart city applications. The journey planner allows a user to receive alerts regarding the events and occurrences (based on her/his chosen topics) in an online form and s/he travels around the city. The smart city dashboard provides an integrated view of the data sources and extracted information from a smart city environment in an effective visual form.

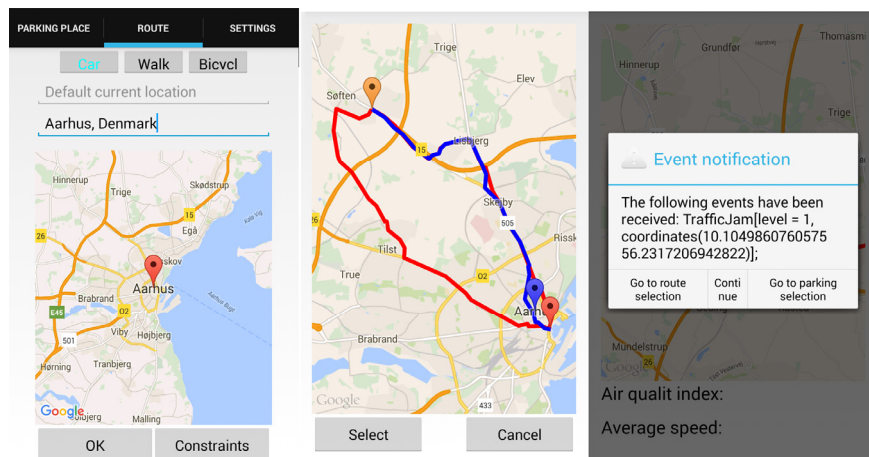


Figure 8: CityPulse Journey Planner Mobile App

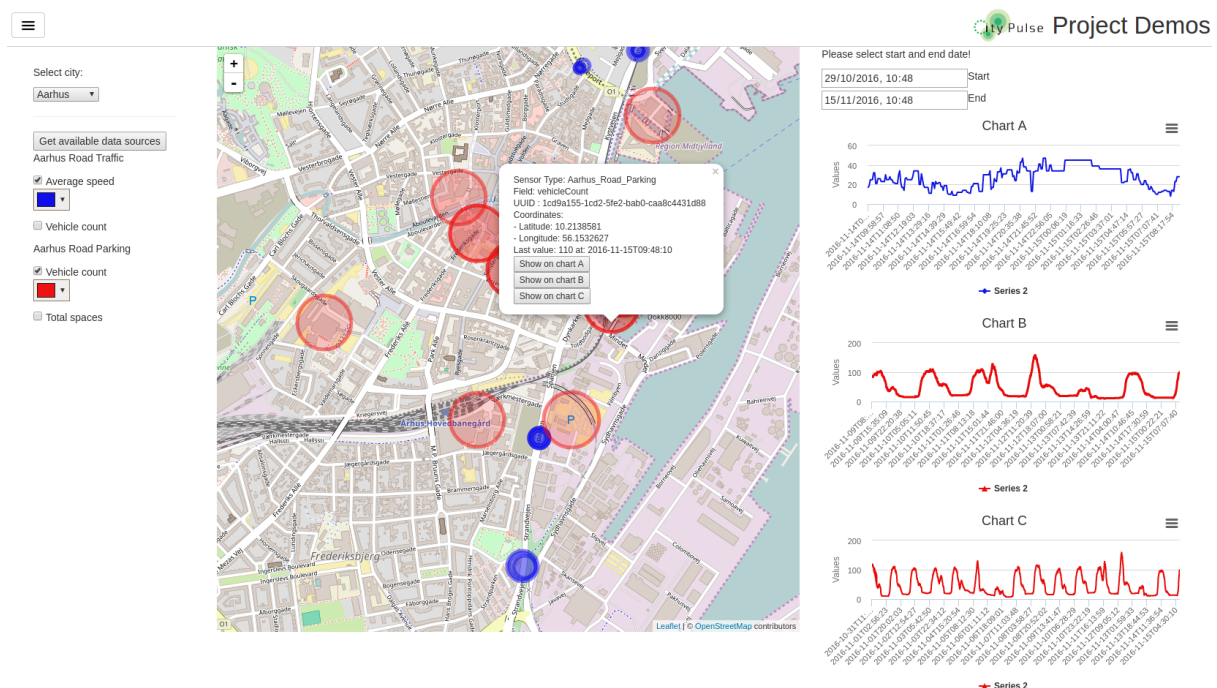


Figure 9: CityPulse Smart City Dashboard

2.2 Scenarios and Smart City Framework

The CityPulse framework supports smart city service creation by means of a distributed system for semantic discovery, data analytics, and interpretation of large-scale real-time Internet of Things data and social media data streams. The framework comprises (1) annotation and aggregation modules adapting automatically to the changes in the input sources in order to minimise information loss, (2) data federation modules finding automatically suitable data input sources at run time according to the user's specifications, (3) complex event processing modules extracting higher level information, (4) quality monitoring modules applying machine learning to assess the quality of the data provided by input sources, (5) contextual filtering modules that constantly monitors the user's current activity to automatically select relevant events, (6) decision support modules which combine semantic technologies and Answer Set Programming to provide an expressive and scalable decision support solution, and (7) powerful visualisation tools. All these modules are provided as open source.

2.2.1 Scenarios and Requirements

The project identified over 101 smart city scenarios and related use-cases in cooperation with partner cities and city cooperation (City Stakeholder Group) and derived a set of requirements for a smart city framework based on proposed use cases, references in the field and "on site" workshops together with city partners. The requirements are encoded in an evaluation metric that has been made available online for the wider community to rank the 101 scenarios. Some of these use-cases were later chosen to be implemented in collaboration with the city partners in the project as the CityPulse demonstrator applications.

The project also provided a method for evaluation of the scenarios based on five higher-level criteria each composed out of several specific criteria each encoding a requirement. Table 1 shows the final metric for evaluation.

User diff. (1)	City relevance (1)	Data streaming (2)	Decision support (2)	Big data (2)
How strong is the expected impact in providing value (e.g. economical, social, etc.)?	Is the scenario culturally relevant?	Is the data accessible (pull/push/subscribe/broadcast)?	How complex is the scenario? (0=simple 3=med 5=high)	Is the data available?
What is the expected uptake?	Is the scenario relevant for citizens?	Is this scenario using a live stream? (Yes/No)	How many data modalities are used? (1=few 3=med 5=high)	Is the scenario scalable?
What is the expected attractiveness and usability?	Is the scenario generally applicable in other cities?	Is there capability in the network to deliver this data stream?	Are there control loops in the scenario? (Yes/No)	What level of privacy consideration does the scenario require?
Is the required data readily and available with the necessary quality and granularity?	Is the scenario relevant for municipalities?	Does the scenario require security (e.g. encryption)?	Is automation included in the scenario? (Yes/No)	
	Does the scenario increase public safety?	Does the scenario require reliability (e.g. data loss)?	Is actuation included in the scenario? (0=no 3=simple 5=complex)	

Table 1: Use-case Selection and Evaluation Matrix

The scenarios are made available online (<http://www.ict-citypulse.eu/scenarios>). The above metrics were made available as an online questionnaire for the wider community to rank the scenarios. The scenarios provide a representative of challenges and potential scenarios for development of smarter city applications. Figure 10 shows the evaluation of the 10 scenarios based on the metrics defined in the project.




Figure 10: Top 10 Ranked Use-case Scenario Evaluations

2.2.2 Smart City Framework

The CityPulse framework contains several components, which can be used for performing real time data analytics on the data generated by the smart cities. Third party developers can use the framework to create new applications for smart cities. Using the framework is simpler and faster than building the application from scratch. Based on the application requirements the developers can deploy the whole framework's stack of components or only a selection of them.

The scope of the framework's common APIs is to allow the developers to configure and use the components based on the processing requirements of the application that have to be developed. The developers can apply the following types of adaptations to the CityPulse framework:

- Register new data sources in the framework;
- Plug into the components custom made processing logic;
- Configure the components operations (e.g. turn on/off a certain feature);
- Configure which components to be used based on the application requirements.

The CityPulse key components are depicted in Figure 11 and can be divided in three main categories:

- Large-scale data stream processing modules: representing the tools, which allow the application developer to interact with the heterogeneous and unreliable data sources from the cities. The tools allow also discovering, summarising and processing the data streams. The components in this category are represented by: resource management, data wrapper, semantic annotation, data aggregation, geospatial data infrastructure, event detection and data federation.
- Reliable information processing modules: these tools can be used to continuously assess the plausibility, correctness and trustworthiness of the data sources. The components from this

category are represented by: atomic monitoring, composite monitoring, fault recovery, conflict resolution and technical adaptation.

- Adaptive decision support modules: containing the tools which can be used for making various recommendations based on the user context and the status of the city. The components from this category are represented by: contextual filtering, decision support and planning, parking, travel tasks.

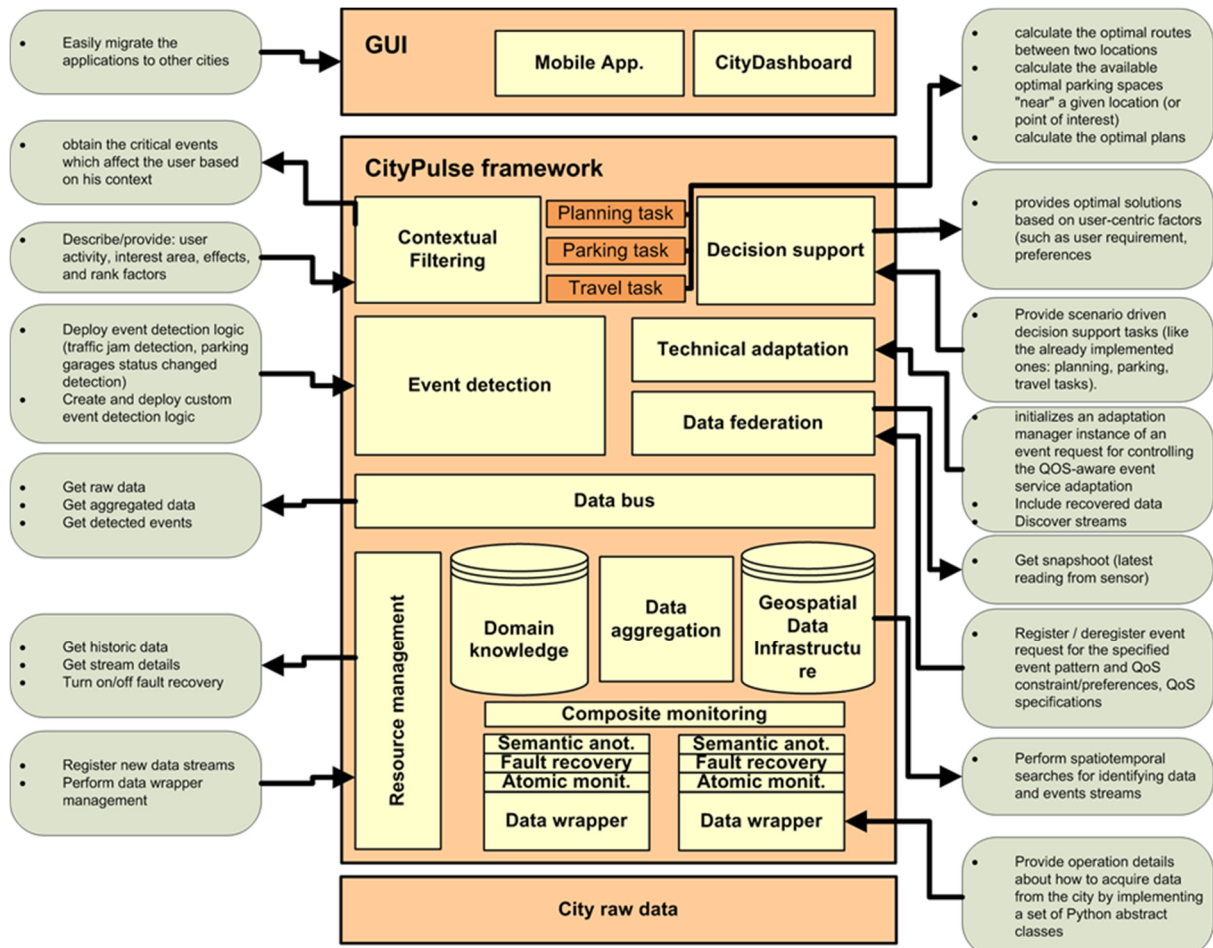


Figure 11: CityPulse Architecture and APIs [Puiu, et. al, 2016]

The CityPulse framework components from the first category fulfil four stages of data processing. First the data wrappers transform the observations that are published by the sensors and provide a unified access layer to the heterogeneous data sources. More precisely, unified access allows modelling the resources in a manner to access these resources systematically in interoperable way regards less of the multi-source, multi-modal smart city data. Next, the semantic annotation component uses semantic descriptions for sensor data, which enables representation, formalisation and enhanced interoperability of sensor data. The data aggregation and summarisation components are used to cope with the large amount of data that has to be processed and stored. In the last stage, the event detection component is used to extract relevant events from the streams of semantically annotated observations.

The components from the second category fulfil mainly a two-staged monitoring approach, with the scope of evaluating the quality of smart city data streams during run-time. Incoming observations are annotated with quality metrics. These Metrics describe the quality level of a data stream and allow consumers (e.g. applications) to evaluate if this stream satisfies the requirements or if the consumer

needs to switch to alternative data streams. Furthermore, the monitoring can be useful to determine the cause of an error. The monitoring component can only be used to find failures and errors in data sources. To solve identified issues additional processing has to be done. Therefore, the CityPulse framework offers additional solutions to deal with faulty data streams. The fault recovery component uses online learning algorithms to estimate wrong or missing values. In case of non-replaceable values, failing data streams, or contradictory information sources, technical adaptation component can be used to automatically switch between available data streams depending on the requirements of the smart city application.

Smart city applications in changing environments require to take into account user preferences and requirements, as well as dynamic contextual information represented by real-time events, in order to provide optimal decision support to the end user at any time. The event-driven adaptation and context-driven user-centricity of the CityPulse framework are materialised by a close loop between the *Contextual Filtering* component, the user application, and the *Decision support* component, which represent the main components from the third category.

The user-centric and event-driven reasoning capabilities of the framework strongly rely on the tight interaction between the user application and the Contextual Filtering component respectively, with the former being in charge of the bidirectional communication with the other two. The user application is also tightly connected with the Decision support component in requesting to compute answers to a decision problem by reflecting constraints and preferences specified by the user himself.

This user-driven loop between the Contextual Filtering and the Decision support component makes it easier for application developer to decide whether to give complete control to the end user on when and how to request adaptation after critical events have been detected, or automatically suggest new solutions.

2.2.3 Reference Datasets for Performance Evaluation

A major issue when testing data set related frameworks is the lack of existing data sets, which are usable for evaluation. CityPulse has collected different datasets during its project lifetime and made them available for download. For an easy re-usage these data sets are annotated with the Stream Annotation Ontology.

The available data sets offer a wide spectrum of different data sources, which are located in the city of Aarhus in Denmark. The underlying data streams are made public by the municipality via an open data portal. This portal (www.odaa.dk) contains static (rarely updated) as well as dynamic (often updated) data sources. As the data portal doesn't store historic data for dynamic data sources the CityPulse project has implemented several data crawlers to get and store the data for a database with historic data sets. Parts of these historic data sets have been annotated with SAO and made available for public download as a dataset collection (<http://iot.ee.surrey.ac.uk:8080/index.html>). In addition to the data sets there is also a collection of tools for handling the data sets (e.g. replaying or validation tools). An additional section for Key Performance Indicators provides ideas on how to measure the performance of frameworks processing smart city data.

By the end of the project the reference datasets include:

- Traffic
- Parking
- Weather
- Cultural events
- Library events

[illegible]

The timespan for the availability of the single data sets varies for the reason of technical or legal issues. Due to server issues some of the data streams were unavailable for a certain time or had to be put out of service for legal reasons (e.g. changed privacy laws). In addition to the open source data sets CityPulse has used some additional data sources, which are not available for free public usage (e.g. TomTom Traffic data) and can therefore not be offered as an annotated data set for public download.

D1.4 Final Report – Dissemination Level: Public

2.3 Large-Scale IoT Stream Processing

The part of the project work developed an open platform for discovery, integration, federation and processing of large-scale IoT streams from heterogeneous sensory resources. The federation enables data stream providers to publish descriptions of their data streams and also provide real-time access to the streaming data. The data streams can consist of various kinds of sensory data provided by machine and human sensory resources. This work package will provide interfaces in addition to open and common description frameworks to publish and describe the stream resources and their data.

2.3.1 Resource Virtualisation and Semantic Annotation

The resource Virtualisation and Semantic Annotation component allows connecting heterogeneous data sources to an integrated framework that describes the various streams while abstracting from the underlying technical details such as the sensor types, the data format and the streaming protocols. Figure 5 showcases the architecture of this component. This enables delivery of large volume of data that can influence the performance of the smart city systems that use IoT data. The semantic annotation uses a lightweight information model that is able to represent the summarisation as well as the reliability and quality of information of both numerical sensor streams and social media data streams (Figure 13).

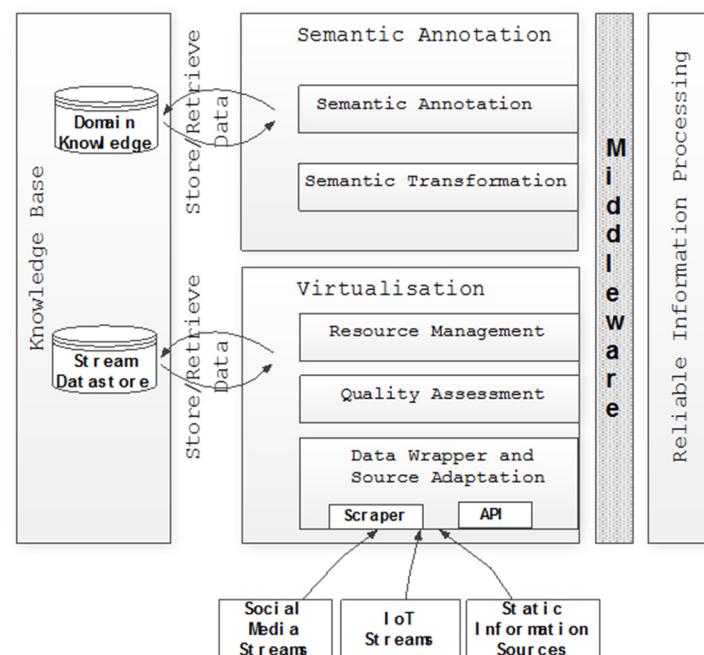


Figure 13: Virtualisation and Real-time Annotation in Smart City Applications

The data wrappers communicate via a middleware with the CityPulse framework. This makes it easy to deploy new data sources can then easily be added in the future just by providing an appropriate data wrapper. The incoming data is semantically annotated using an ontology that can represent the domain knowledge in order to deal with various forms of heterogeneity, including data modality, data model (e.g. terminology), and data representation (e.g. RDF/JSON). Utilisation of semantic technologies for IoT enables the interoperability among various IoT resources, data providers and consumers.

2.3.2 Data Federation and Aggregation

The data federation and aggregation component provides the next step to extract actionable information from the data streams. Data federation is a core component for the smart city framework, as it is responsible to process application request for IoT streams and automatically discover the most relevant

data streams after catering individual requirements and preferences for a particular user request. Data federation is also responsible for automatically integrating heterogeneous data streams and performing complex event processing over the integrated data streams. In order to address the aforementioned tasks for the data federation component, we propose to integrate Semantic Web (SW), Service Oriented Computing (SOC) and Complex Event Processing (CEP) and provide a data federation solution for smart city applications based on event services. We refer to an event service providing complex events which describes the complex event pattern in its service description as a Complex Event Service (CES), otherwise if the event service provides primitive events or do not describe event patterns in its service description we call it a Primitive Event Service (PES).

A service network consisting of CESs and PESs is called an Event Service Network (ESN). We refer to an event service that delivers semantically annotated events and describes the service metadata with semantic annotations as a Semantic Event Service (SES). Figure 14 shows the Architecture of the data federation and aggregation component and how it interacts with the data virtualization and semantic annotation component.

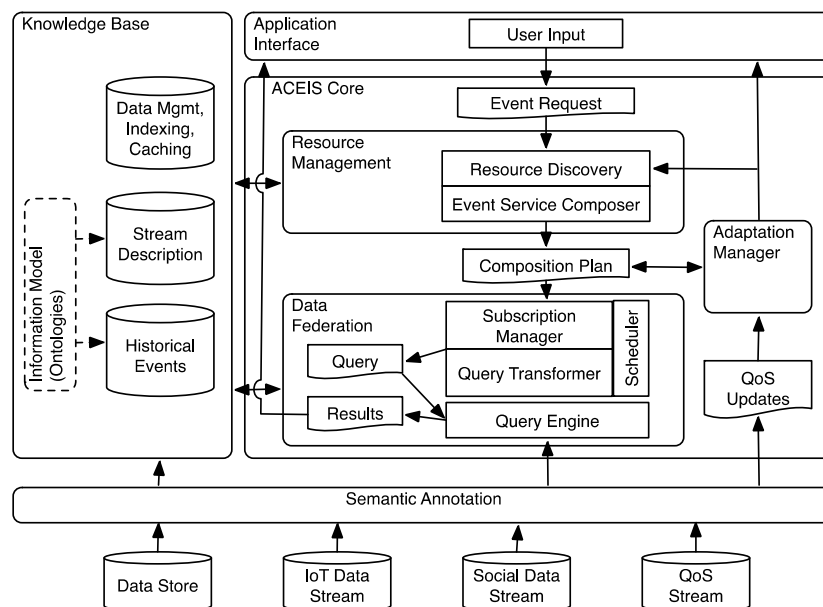


Figure 14: Architecture of the Data Federation and Aggregation Component

2.3.3 Event Detection for Urban Data Streams

The stream Event detection component provides the generic tools for processing the annotated as well as aggregated data streams to obtain events occurring into the city. This component has to be highly flexible in deploying new event detection mechanisms, since different smart city applications require different events to be detected from the same data sources [Puiu, et al., 2016]. The component has been developed using the Esper engine.

Usually, the development of an event driven application consist of two main steps: 1) real-time data acquisition, interpretation and validation 2) execution of event detection mechanism in order to detect the patterns of events.

The first step is performed automatically in the CityPulse framework. When a new specific event detection mechanism is deployed for processing the data coming from a set of streams, the Event detection component performs automatically (with no intervention from the application developer) the following actions:

- Makes a request to the Resource management to identify the description of the streams (i.e., the routing keys where the requested streams are published on the Data Bus, and the details about how to interpret an observation);
- Connects to the Data Bus and continuously converts the received observations from RDF format to the one requested by the Esper engine.

In order to fulfil the second step of the event detection mechanism the application developer has to develop an event detection node, under the form of a Java class, which contains the event detection pattern. The event detection node (Figure 15) can be seen as a black box with inputs being input streams from Data wrappers (DWS) and configuration parameters, and having as output the stream of detected events (OES).

In order to develop a new event detection node, the application developer has to extend a dedicated Java class, which provides methods for defining the event detection pattern. Existing event detection nodes can be reused by simply changing the configuration parameters and the input streams.




Figure 15: Event Detection Node

2.3.4 Knowledge Acquisition Toolkit (KAT)

The Knowledge Acquisition Toolkit provides capabilities for analysing data. Various methods have been implemented for data pre-processing to bring the data into shape for further processing, dimensionality reduction to either compress the data or reduce its feature vectors, feature extraction to find low-level abstractions in local sensor data, abstraction from raw data to higher-level abstractions and finally semantic representations to make the abstracted data available for the end-user and/or machines that interpret the data. It allows the user to see how different methods combined with each other can be used to extract information from their data set. This helps to pick the right methodologies used for each application specific task. The gained insights can then be used to select the right methods and configure the data analytics for each data wrapper in the resource management component. A screenshot of the toolkit is shown in Figure 16. The toolkit is available as an open-source component at: (<https://github.com/CityPulse/Knowledge-Acquisition-Toolkit-2.0>).

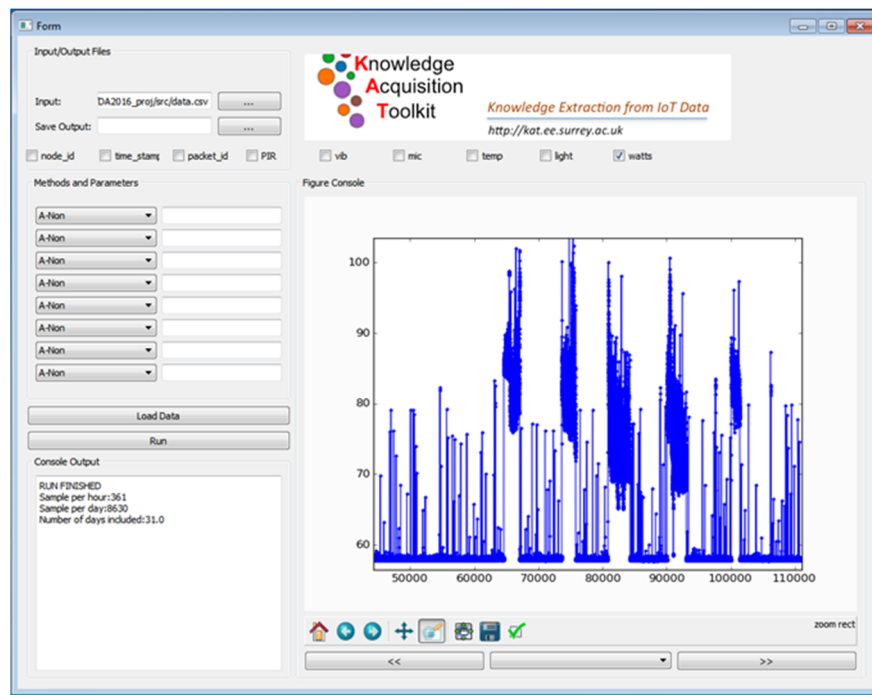


Figure 16: KAT Data Analytics Toolkit

2.3.5 Resource Management

The Resource Management component is responsible for managing all Data Wrappers. During runtime an application developer or the CityPulse framework operator can deploy new Data Wrappers to include data from new data streams. The folder "wrapper_dev" contains exemplary Data Wrappers for traffic and parking data of the city of Aarhus, Denmark as well as weather, air quality and incidents of the city of Brasov, Romania. The Resource Management component can be used for the following types of scenarios:

- Fetch live stream data via one or more Data wrappers
- Replay historic data embedded in a Data wrapper

The Resource management uses a configuration file to store the connection details to other components of the framework. The configuration is provided as JSON document named "config.json" within the "virtualisation" folder. The configuration consists of a dictionary object, where each inner element holds the connection details to one component of the framework, also as a dictionary. The resource manager is available via the CityPulse GitHub code repository at: (<https://github.com/CityPulse/Resource-Manager>).

2.4 Reliable Information Processing

The results achieved for reliable information Processing show that, depending on the availability of individual data sources, various algorithms and processes can be used to evaluate city infrastructure data and infer decisions on data source usage. Therefore, spatio-temporal relations between events and continuous measurements are used to determine plausibility of the evaluated datasets. CityPulse ensures reliable information processing through four measures, namely Testing, Monitoring (Atomic and Composite), Fault Recovery, and Conflict Resolution, which are iteratively described in this deliverable. Furthermore, tools to visualise and evaluate the performance of these measures are presented.

The following Figure 17 provides an overview about all the components realised within the Reliable Information Processing components and their relation to other components. They are namely: Quality Ontology, Atomic and Composite Monitoring, Fault Recovery and Testing.

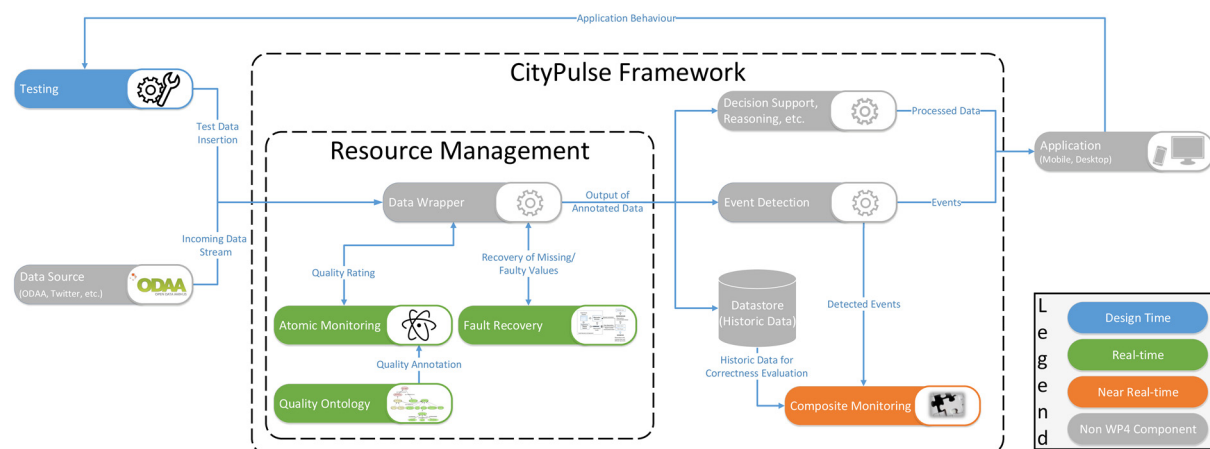


Figure 17: Reliable Information Processing Overview

2.4.1 Accuracy, Trustworthiness and QoI

One major key point of the CityPulse project is the Reliable Information Processing. This ensures that other components of the framework can operate within defined quality limits of the received data from various heterogeneous data streams. To store the calculated quality of received data an information model has been developed.

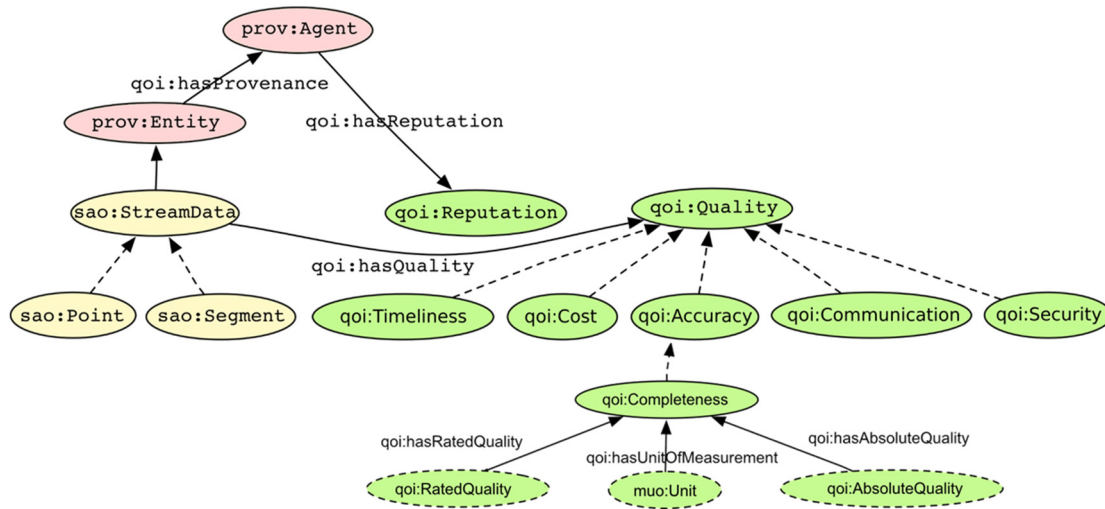


Figure 18: Quality Ontology

Within the CityPulse framework processed information is semantically annotated. For this reason, we developed an ontology to connect the calculated Quality of Information directly with the information itself. The developed Quality Ontology¹ makes use of the existing PROV-O ontology and the Stream Annotation Ontology (SAO) developed within CityPulse [Iggena et al., 2014]. Figure 20 depicts the structure of the ontology. The red coloured parts are imported from PROV-O whereas the yellow classes are part of SAO. Via the *hasQuality* property it is possible to annotate every data annotated with SAO with QoI. Our information model for this *QoI* contains five categories each having some subcategories. In the provided figure one example, Completeness as a subcategory of Accuracy, is depicted. The subcategory itself has three properties. One contains an absolute quality value (e.g. 90% for Completeness), the second one specifies the unit of the provided absolute value whereas the rated quality holds a value calculated with a Reward and Punishment algorithm.

2.4.2 Conflict Resolution and Fault Recovery

The Fault recovery component ensures continuous and adequate operation of the CityPulse application by generating estimated values for the data stream when the quality drops or it has temporally missing observations [Puiu et al., 2016]. When the quality of the data stream is low for a longer time period, an alternative data source has to be selected. The selection can be performed automatically by the Technical adaptation component. In other words, the technical adaptation process does not have to be triggered if the QoI of a stream is low only for a short period of time because the Fault recovery component provides estimated values.

The Fault recovery component is integrated into the data wrapper. The fault recovery mechanism is triggered to generate an estimated value when the atomic monitoring component has determined that the current observation is invalid or missing.

The building blocks of the Fault recovery component are presented in Figure 19. The component has a buffer which temporary stores the latest observations generated by the data stream, and a reference dataset that contains sequences of valid consecutive observations from the data stream. When an estimated value is requested, the k-nearest neighbour algorithm [Richeard et al., 2000] is used to select

¹ <http://purl.oclc.org/NET/UASO/qoi>

a few sequences of observations from the references dataset, which are similar to the current situation. At the end the estimated value is computed from the selected sequences of observations.

During the normal operation, when the QoI of the stream is high, the fault recovery component extends the reference data set with the sequences of observations from the buffer if a similar signal pattern was not included before.

Initially, when the Data wrapper is deployed, the reference data set is empty and based on the normal operation (from stream quality point of view) it is extended. As a result of that, the work of the 3rd party application developer is reduced, because he does not have to collect historic data from the stream, to clean and to validate it in order to create the reference dataset. Using the API exposed by the resource management, the 3rd party application developer can turn on and off this component based on the CityPulse application requirements.

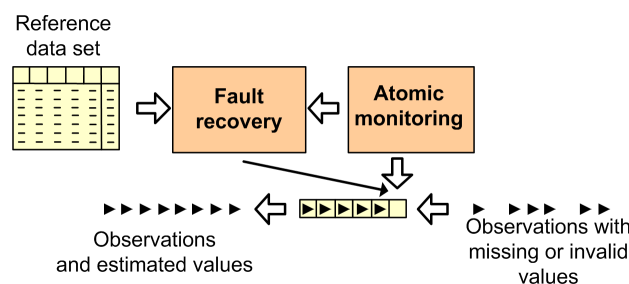


Figure 19: Fault Recovery Component Workflow

2.4.3 Monitoring and Testing

The availability of smart city applications highly depends on the availability of appropriate, accurate, and trustworthy data. This includes the availability of necessary data sources as well as accessing their QoI descriptions. The reliability of the extracted sensor information has to be monitored during run-time. Monitoring methods are used to compare the information quality of data streams with the QoI-requirements of an application. To get a continuous view of sensor stream qualities two different monitoring components have been designed. The Atomic Monitoring component focuses on sensor information of an individual sensor stream whereas the Composite Monitoring component analyses inter-dependencies between various sensor streams to evaluate the plausibility of information.

The separation into two components ensures a constant real-time quality annotation for basic quality parameters and a more complex proof of correctness to rate the trustworthiness of data sources, as it probably could not to be done in real-time.

2.4.3.1 Atomic Monitoring

The Quality Ontology is used by the Atomic Monitoring component which is responsible to annotate incoming data with QoI. Within the CityPulse framework this is achieved by a direct integration of the Monitoring into the data wrapper, which are receiving information from incoming data streams. This direct integration has been implemented to fulfil the real-time requirements of the framework. More complex functionalities are done in the Composite Monitoring, which is decoupled from the Data Wrappers (see section 2.3).




Figure 20: Atomic Monitoring Functionality

The basic principle of the Atomic Monitoring is shown in Figure 22. Within the CityPulse framework each data stream has a description, which contains attributes of the stream and the provided data (e.g. update interval, provided data fields). For an incoming observation the Atomic Monitoring compares the description with the received data. Based on this comparison the rated QoI for the stream is increased or decreased with the help of a reward and punishment algorithm. The calculated QoI is returned to the overlaying Data Wrapper and annotated with the Quality Ontology.

For the Atomic Monitoring five Quality Metrics have been developed: Age, Completeness, Correctness, Frequency, and Latency. The following example illustrates the processing of the metric Frequency. For an incoming observation or after a certain time (update interval) the time between the last update is compared with the planned update interval stated within the stream description and with the result the QoI is rated (see Figure 21).




Figure 21: Determination of the QoI Metric Frequency

For implementation and evaluation purposes, the CityPulse team has developed a data emulator that re-plays the traffic data stream for the city of Aarhus that is collected in December 2015. This data stream contains 449 single traffic sensors, whose QoI has been calculated. The following graph shows the cumulative distribution of the quality level for these 449 sensors.

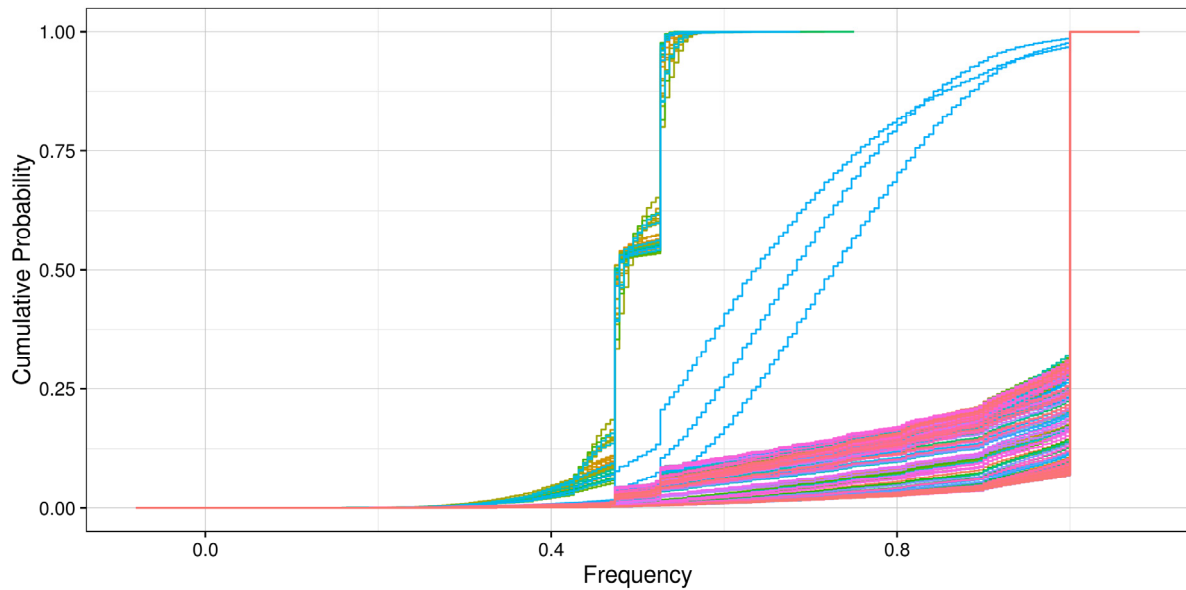


Figure 22: Update intervals for the traffic data in Aarhus

Most of the sensors fulfil the update interval in more than 80% of the time (see Figure 22). However, a number of sensors are fulfilling the Frequency metric for only half of the time. At the initial stage we were not able to find the reason for this; however, after further discussion the city of Aarhus team and also further evaluation of the analysis results, we noticed that the traffic sensor data was collected by resources that were deployed on lampposts and in some cases (especially in summer time) in which the lampposts did not have power supply the sensors were not fully charge and this led to loss of some parts of the sensory data over time.

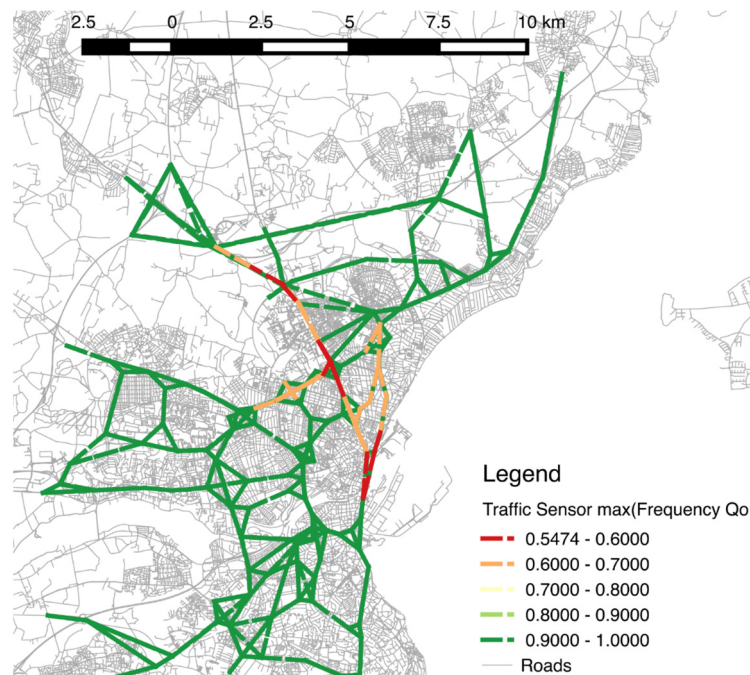


Figure 23: Aarhus Data Frequency Map

2.4.3.2 Composite Monitoring

Figure 24 describes the evaluation process of the Composite Monitoring. The goal is to determine correctness value (C_e) for event (e). A set of correlating data streams (S_e) is used as validation source. The Figure shows the 4 phases of the validation process:



Page 33




Figure 25: Example of Space-Time Resolution of Sensors and Events

The overall results of the available datasets show that for more than 75% of the severity level 4 (Closed Road) the raw data of a nearby sensor measured a vehicle count of 0 cars during the event period. Figure 26 shows the cumulative density function (CDF) of the irregular incidents. It shows that for reported events most vehicle speed readings are slower than described by the Seasonal Component S_t , which means that the Irregular Component $I_t < 0$. The change of the vehicle count does not clearly reflect the event. The vehicle count can drop because of a traffic jam or a traffic jam can be the result of the road being overloaded.




Figure 26: Avg. Speed (Irregular) and Vehicle Count (Raw) Change During Traffic Incidents

Furthermore, also continuous sensor readings have been used to monitor the data stream reliability and to illustrate the component validity [Kuemper et al., 2016]. Commercial TomTom Traffic Flow data has been evaluated against ODAA Traffic Flow measurements. For every TomTom measurement that showed a severe slowdown of more than 15km / h compared to the free flow speed an evaluation has been triggered. Figure 27 shows the heat map distribution of 28306 pairwise comparisons of the two data sources. The colour in the heat map shows the density of measurement points. The figure depicts that there is no clear linear correlation between the two data sources. 95.5% of ODAA comparisons also have an irregular component value $I_i \leq 0 \text{ km / h}$, which is indicating a decrease in speed. These measurements confirm the plausibility of the TomTom Flow data, although there is no direct correlation between measured values.

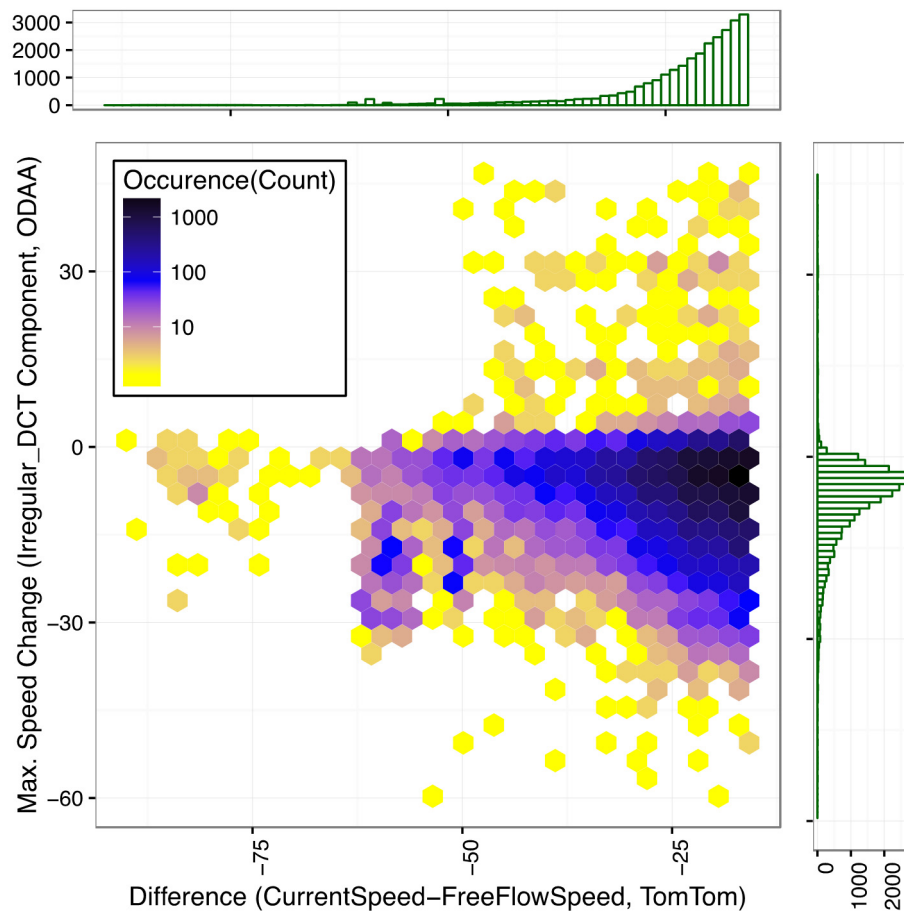


Figure 27: Heat Map of Measured Value Pairs of Measured Traffic Speed Changes

Overall, the above discussed efforts in the project provided various measures for ensuring and increasing the reliability of smart city applications by employing a monitoring approach for exemplary correlating data streams. The monitoring can be triggered by detected incidents or suspicious sensor readings in data streams. It provides a correlation dependency model-based approach integrating new evaluation schemes. This approach supports appropriate spatiotemporal distance measures to achieve an efficient monitoring of relevant correlating data. The suggested framework provides methods to cope plausibility analysis for heterogeneous data sources in smart city applications, and in addition considers that the sensors can become unreliable over time by approaching a continuous time series evaluation. As a countermeasure, several actions to identify and react on changes in the quality of information have been investigated and integrated into the CityPulse framework. In the future, we plan to investigate an open platform for a drag and drop evaluation approach, which allows applying the concept to different domains and allows easy utilisation of city infrastructure knowledge.

For scalability reasons and to be able to meet real-time requirements the monitoring is divided into a two-stage approach. The first stage, Atomic Monitoring, is responsible for single stream data quality metrics, such as Completeness, Age and Frequency. To achieve best possible performance, the Atomic Monitoring was integrated directly into the data streams Data Wrappers, a modular concept acting as entry point for external data resources into the CityPulse framework. The second stage, the Composite Monitoring, triggered only by detected events, rapidly dropping QoI values or through manual interaction, provides a multi information source plausibility evaluation scheme. For this correlating information sources within a spatio-temporal distance around the event are identified. The work here

highlights the importance of a fitting distance model describing to propagation properties of a physical phenomenon along with an appropriate correlation method.

2.4.3.3 Testing

The goal in CityPulse is to test the reliability of smart city applications with respect to the reliability of the external resources required for its execution. For this a series of test cases has to be generated, where each test case 'simulates' a less reliable deployment of external resources as the former. Since a true ground truth is not accessible, the first test case (TC₀) in such a series uses unmodified historic data, which are represented by the Reference Dataset from the CityPulse project, and acts as ground truth. In each following test case the output of the CityPulse framework is recorded. The distances between the inputs of a test case to TC₀ and the distances of the outputs of same test cases must correlate or lie below a threshold in order to pass a test case.

The execution of a test case utilises the replay mode of the Resource Management, which is already capable of using historic datasets instead of live data. A setup for a test case is therefore the replacement of the original historic datasets with the manipulated, degenerated datasets. To limit the number of possible external resources involved in a test case an application profile states the types of resources it uses during execution. A geo-spatial search for relevant external resources for a specific test scenario will further reduce the test case input space. Both measure reduce the test execution time significantly. Figure 28 illustrates the test execution process and highlights the loop, in which the ground truth is degenerated up until the application output changes sufficiently to be able to make a statement about its robustness. For the sake of simplicity, the loop's end condition in case the maximum number of iterations was executed is not shown in the figure. The next section will describe in more detail the degeneration process of the historical data. In contrast to monitoring the testing is executed during the design-time of an application respectively before the deployment of an application.

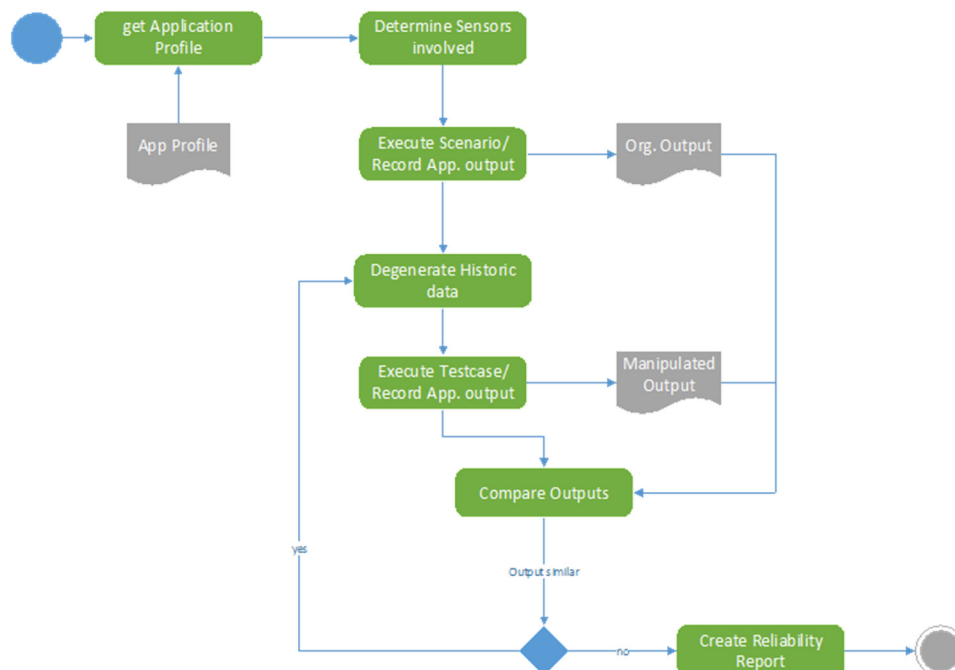


Figure 28: Execution of Test Cases

Testing deals with the stability of smart city applications with a decreasing quality of sensory observations from deployed sensors. Initially, we describe the effects of faulty sensors on the output signal of measurement equipment. This information is then used to define an iterative test case

generation process, where each successive test case stimulates the CityPulse framework with less reliable data streams (lower quality) than its predecessor. The process is repeated until the framework's output differs sufficiently from the original output (test run using original historic data) in order to make a statement about the quality threshold required by the application under test.

2.4.4 QoI Explorer Open Source Tool

The Atomic Monitoring component is directly integrated in the Wrappers of the CityPulse frameworks and has no direct user interface. To enable domain experts to inspect the current QoI levels of the integrated data streams a tool called Data Quality Explorer has been developed. The Explorer is a web based tool, which offers the possibility to inspect the QoI values of the data streams and their sensor. The single sensors are shown on a map providing a colour based quality view. With additional statistics for all implemented QoI metrics it is possible to find malfunctioning sensors. In addition to the QoI view we have integrated a time based performance measurement for the CityPulse framework showing the processing times of the single components. In addition, the tool provides an offline mode for demonstration purposes. A sample screenshot of the main view with the map showing the city of Aarhus and the deployed sensors is shown in Figure 7. For more information, there is a detailed video for all the features on YouTube (<https://www.youtube.com/watch?v=edFrlciO8Og>) and the source code is available at GitHub (<https://github.com/CityPulse/Data-Quality-Explorer>).

2.4.5 Event Test App

To test the events within the Composite Monitoring an Android application to inspect, create, delete, and change events has been implemented. Figure 29 shows a screenshot of the application. The application is directly connected to the CityPulse framework by subscribing to the events sent on the message bus.



Figure 29: Event Report Mobile Test App

Events generated by the integrated Event Detection of the CityPulse framework are shown as well as events reported by other users. For user-generated events it is possible to change their intensity or to delete them. The applications source code is available at: (<https://github.com/CityPulse/EventReportApp>).

2.5 Real-Time IoT Intelligence

The key accomplishment of the Real-time IoT Intelligence is the provision of an intelligent layer which can dynamically process user/application requests and provide decision support system by performing adaptive complex reasoning over streaming data. Real-time IoT Intelligence can correlate user context and contextual events to refine those critical contextual events which can potentially effect plausibility of the extracted results. In adaptive urban reasoning, the Technical Adaptation component monitors quality of IoT streams and adapts to the next available stream if quality of the existing streams is declined, while Contextual Filtering component evaluates criticality of any given events to determine if the event can have any potential effect on user's context. For user-centre decision support in dynamic environments, a rule-based Decision Support component processes application request containing users' constraints and preferences and calculates the best answers satisfying user and application defined requirements and preferences.

The Figure 30 provides an overview of the main components developed in Real-time IoT Intelligence layer and presents the information flow within its components as well as their interactions with other related components of external work packages.

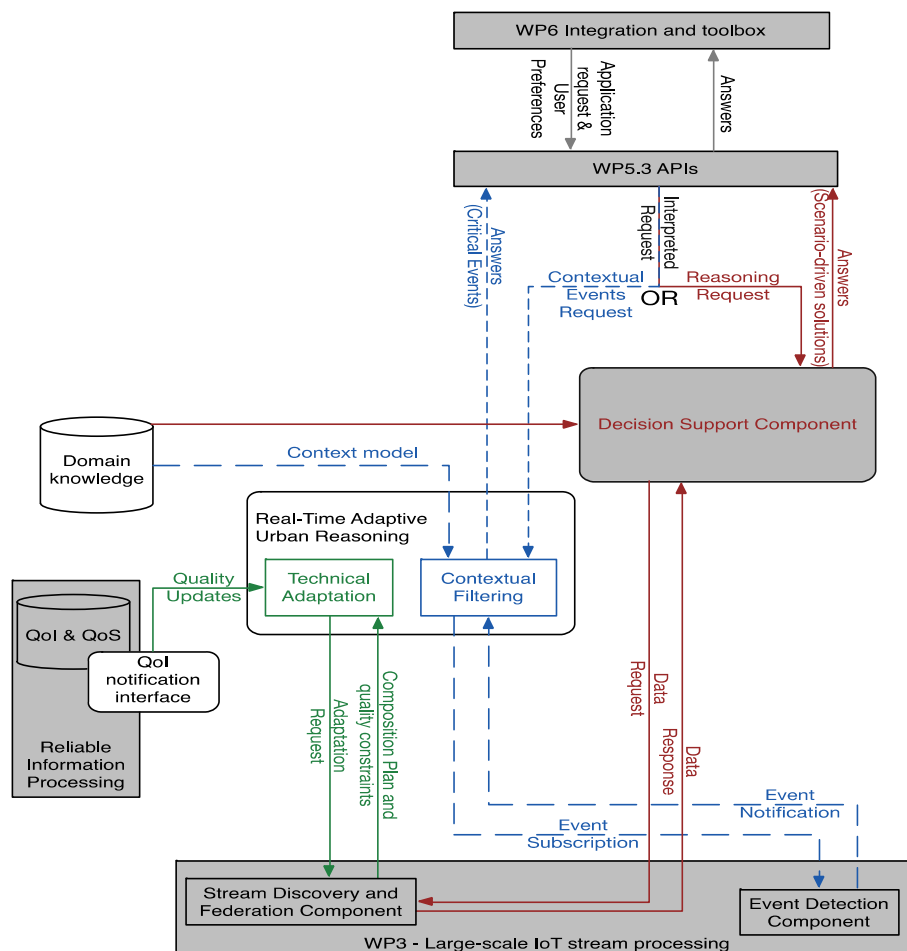


Figure 30: Overview of the Real-time IoT Intelligence Components

Three main functionalities in the CityPulse Framework are key for Real-Time IoT Intelligence:

- Knowledge-based **adaptivity** relying on continuous identification of what events are relevant and critical at any given time for a particular user in a particular application
- High-level **reasoning** to extract meaningful information for contextual and user-centred decision support and problem solving
- Dynamic and intuitive Visualization

In addition to that, relying on the data federation components, the reasoning capabilities of the CityPulse framework can achieve dynamic selection and composition of the best relevant streams and events at any given time.

In the following three subsections, we discuss the results achieved for Real-time IoT Intelligence in terms of contextual adaptivity and user-centred reasoning for decision support, and we illustrate the key role of Visualization to provide access to the results of IoT-Intelligence processing at different levels that go beyond the end-user application.

2.5.1 Real-Time Adaptive Urban Reasoning

Adaptive Urban Reasoning in CityPulse framework is mainly concerned with providing a higher-level control to (i) configure data streaming in terms of discovery, federation and compositions of event streams, and (ii) identify unexpected events that are computationally relevant for the decision support functionality of the CityPulse framework. A key aspect in adaptive urban reasoning is to define the mechanism to identify users' requirements, constraints and preferences [Gao et al., 2015]. The CityPulse framework caters for user defined requirements and collects the following information about the users' requirements:

- Functional Details: represent the qualitative criteria used in the reasoning task to produce a solution that best fits the user needs.
- Functional Parameters, defining mandatory information for the Reasoning Request (such as start and end location in a travel planner scenario).
- Functional Constraints, defining a numerical threshold for specific functional aspects of the Reasoning Request (such as cost of a trip, distance or travel time in a travel planner scenario). These restrictions are evaluated as hard constraints which need to be fulfilled by each of the alternative solutions offered to the user.
- Functional Preferences, which encode two types of soft constraints: a qualitative optimisation statement defined on the same functional aspects used in the Functional Constraint (such as minimisation of the travel time); or a qualitative partial order over such optimisation statements (such as preference on the minimisation of the distance over minimisation of the travel time). Preferences are used by the Decision Support component to provide to the user the optimal solution among those verifying the functional constraints.

2.5.1.1 Technical Adaptation in Stream Processing

In CityPulse framework, the Technical Adaptation component provides an adaptive loop for IoT streams discovery and adaptation. This component is responsible to continuously monitor the QoI/QoS scores of the all the contributing data streams relevant to a specific event request. This component comprises of three main modules, namely:

- Discovery: This module receives a composition plan with a reference to the event request for which the composition plan was generated. Event request contains user/application functional requirements as well as non-functional requirements (QoI/QoS constraints and preferences). Discovery module discovers the relevant QoI/QoS update streams to subscribe.

- **Monitoring:** This module generates a subscription request for any update in the QoI/QoS scores of the relevant data streams. Monitoring module continuously monitors and verifies that whether QoI/QoS scores of all the contributing data streams in a composition plan are compliant to the event request.
- **Adaptation Request:** This module is triggered if any of the user defined non-functional constraints and preferences are violated, this module requests the federation component to re-discover the data streams, which qualify to the user constraints and preference, described in the original event request. Federation component will send the new composition plan to the Technical Adaptation, which will resume discovery and monitoring module for the new composition plan. However, if the Federation component is unable to re-generate new composition plan because none of the available data stream comply with the constraints and preferences defined in the event request. The Technical Adaptation component provides a feedback to the user/application to reconsider the constraints/preferences and generate a new event request [Gao et. al, 2016].

Due to the optimisation issues and the close integration of Federation components with Technical Adaptation, CityPulse framework bundled the Technical Adaptation inside the Federation component. Figure 14 represents technical adaptation components integrated within the architecture of data federation and aggregation component. The source code of Technical Adaptation component is integrated within Federation component and available at CityPulse GitHub (<https://github.com/CityPulse/Stream-Discovery-and-Integration-Middleware>).

2.5.1.2 Contextual Filtering

In Contextual Filtering, the adaptive behaviour realised in the CityPulse framework makes it possible to identify which unexpected events could affect the optimal result of the user-centric decision making. Based on these events, a feedback is provided to the reasoning component that triggers the computation of new solutions whenever the problem solution is compromised.

When an application requires a specific decision support solution to be computed, the adaptive mechanism implemented in the Contextual Filtering component subscribes to classes of events that are relevant for the particular application (as specified by the application developer through semantic event description available from the Events-Context Ontology) and compute their criticality to the specific problem being solved. The use of semantics makes this process interoperable, traceable, and flexible across several applications.

The conceptual interactions between the Contextual Filtering component and the other functional components in related technical work packages is illustrated in Figure 31. The Contextual Filtering component gets as input; i) an answer to event requests as an event stream returned from event detection component and ii) contextual Information from the decision support component (including user query, query response, user dynamic context and user feedback when available).

The output provided by the contextual filtering component consists of a list of critical events that affect the query response, which is passed on to the decision support component to take appropriate actions.




Figure 31: Contextual Filtering Interactions

2.5.2 User-centric Decision Support in Dynamic Environments

The Decision Support component of the CityPulse framework represents higher-level intelligence, and the main role of this component is to enable reactive decision support and problem solving functionalities to be easily deployed, providing the most suitable answers to a specific user at any time.




Figure 32: Event-based User-Centric Decision Support

The user-centric and event-driven features of the Decision Support component rely on the tight interaction with the user application and the Contextual Filtering component respectively, with the former being in charge of the bidirectional communication with the other two, as represented in Figure 32.

As illustrated in Figure 33, along with the reasoning request represented by the set of rules generated from the user request, the Decision Support component also receives input from background knowledge containing static information (such as maps), the incoming flow of events it subscribes to, and the result of partial answers from external modules.

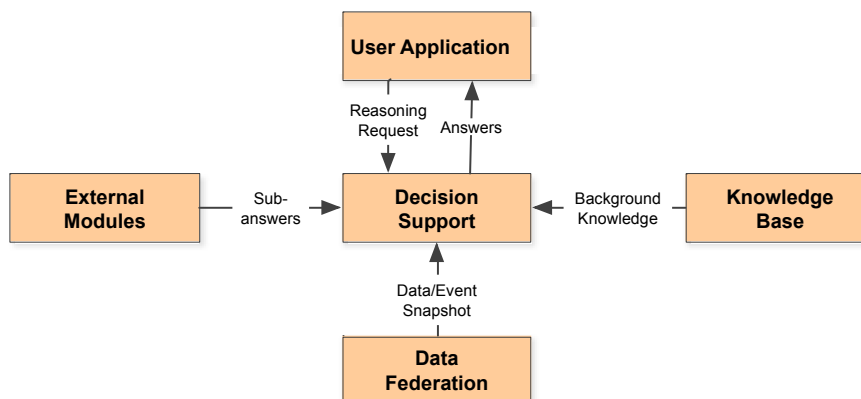


Figure 33: Decision Support Interactions

The reasoning capabilities needed in CityPulse to support users in making better decisions require handling incomplete, diverse and unreliable input, as well as constraints and preferences in the deduction processes. Semantic technologies for handling data streams cannot exhibit such complex reasoning capabilities. Conversely, state-of-the-art logic-based non-monotonic reasoners can perform such tasks but are only suitable for data that changes in low volumes at low frequency. In CityPulse, we bridge this gap by enabling complex reasoning to be performed in a scalable way over data streams [Mileo et. al., 2015]. While the expressivity in the Decision support component is achieved by using a declarative non-monotonic logic reasoning approach based on Answer Set Programming, scalability is

achieved by i) using problem decomposition via external modules that are used to compute partial solutions used by the Decision Support component for contextualisation and optimisation, ii) using continuous query processing and event filtering to reduce the size of the input, and iii) through parallelisation over subsets of the input.

Problem decomposition refers to the use of specific computational tools that can help finding partial solutions faster. As an example, for optimising user-centric and adaptive routing task, the Decision Support component relies on the functionalities of a Geo Spatial database to quickly limit the set of potential candidate routes in a limited geographical area. Since Answer Set Programming computational bottleneck is the grounding phase, where rules with free variables are instantiated with all possible combination of input atom, the size of the input has a heavy effect on scalability. In CityPulse we address this problem not only by reducing the size of the input via semantic query processing and contextual event filtering, but also by designing mechanisms for parallelisation over input partitions.

When parallelising the execution of logic program, state-of-the-art techniques focus on partitioning the rule-set, computing partial results and then joining partial results for the final output. Given that most of the decision support problems in the CityPulse scenario are not stratified (e.g. rules are all dependent on each other and cannot be split), in our solution we focus on input dependency, which can be observed more often and identified more efficiently than rule-dependency. We investigate how executing parallel instances of the same logic program, each on a subset of the input stream, can also have a positive effect on scalability. In our preliminary experiments (under submission), we define the notion of “input dependency graph” and split it with duplication guided by heuristics. Each split is run over parallel instances of the logic program. Despite further optimisation are possible, our initial experiments demonstrate that partitioning based on the input dependency graph reduces the latency by 50%, still maintaining the accuracy of the results, paving the way for scalable approaches to complex reasoning over data streams.

2.5.3 Visualization in Dynamic Smart City Environments

The CityPulse framework exposes the following main user interfaces:

- CityPulse dashboard: supports visual analytics for the data sources registered in the framework;
- CityPulse QoI explorer: GUI to monitor in real-time the quality and the trust of the data sources;
- CityPulse components performance monitoring: GUI which can be used to monitor in real-time the performance (processing times) of various CityPulse components;
- CityPulse social media analysis: a GUI which displays the events extracted from the social media (Twitter) streams;
- CityPulse 3D map: support 3D visual analytics for the data sources registered in the framework.

The CityPulse dashboard provides immediate and intuitive visual access to the results of its intelligent processing and manipulation of data and events. The ability to record and store historical (cleaned and summarised) data for post-processing makes it possible to analyse the status of the city not only on the go but also at any point in time. In addition to these, the dashboard enables the diagnosing and “post mortem” analysis of any incidents or relevant situation that might have occurred. To facilitate that, a dashboard for visualising the dynamic data of smart cities is provided on top of the CityPulse framework. Based on this dashboard, the user has the possibility to visualise a holistic and summarised view of data across multiple contexts or a detailed view of data of interest, as well as to monitor the city life as it evolves and as things happens. The investigation of past city events or incidents can be conducted from different perspectives, e.g. by observing the correlations between various streams, since the streaming data is stored in the framework for a period of time which can be configured, and it can be retrieved for visualisation and analysis at any moment.

To allow the domain expert to observe the quality of data streams the CityPulse framework provides a tool called QoI explorer to monitor deployed sensors in the city. Within this tool, a map visualises the location of each sensor. By selecting a quality metric in the menu at the right side the state for this metric is shown by a colour for each sensor on the current map. A legend in the left bottom corner indicates the corresponding quality value for the sensors colour. By clicking a sensor on the map a detailed view for the selected sensor is shown. In addition, the menu at the right side provides a selection of the time range for shown quality values and to switch between average, minimum, or maximum values. A histogram for each metric allows a quick overview about the distribution of quality values for the sensors currently shown within the map.

The QoI Explorer offers an additional tool to measure the current performance of the CityPulse framework. All components, which are integrated into the Resource Management can be monitored for their current average time consumption.

The 3D map tool provides an innovative way of visualising real-time data related to a city in a 3D model. By using this tool municipality employees, city planners or citizens in general, are able to have a perception of different data and its effect on the 3D model. This way users are also able to identify data patterns and other phenomena in the different areas of the city.

2.5.4 Decision Support and Contextual Filtering Open Source Tool

The Decision Support and the Contextual Filtering are implemented in the same package (as Java JAR libraries) as two different web-socket endpoints of the same server. The Decision Support component is automatically deployed when the Contextual Filtering component is deployed and vice versa.

The main role of the Decision Support component is to enable reactive decision support functionalities to be easily deployed, providing the most suitable answers at the right time by utilising contextual information available as background knowledge, user patterns and preferences from the application as well as real-time events. The reasoning capabilities needed to support users in making better decisions require handling incomplete, diverse input, as well as constraints and preferences in the deduction process. This expressivity in the Decision Support component is achieved by using a declarative non-monotonic logic reasoning approach based on Answer Set Programming. The Decision Support component produces a set of answers to the reasoning request that satisfy all user's requirements and preferences in the best possible way. These solutions are computed by applying sets of rules deployed as scenario-driven decision support modules. We currently support three different types of decision support modules, covering a broad range of application scenarios:

- Scenario 1 (Routing APIs): It provides the best solution(s) for a routing task. The users not only identify starting point and ending point but also are able to provide various selection criteria (including constraints and preferences) in order to select the optimal routes.
- Scenario 2 (Parking space APIs): It provides the best selection among a set of available parking places based on optimisation criteria, constraints and preferences of the users.

The components and their technical documentation are available as open-source via the project's code repository at: (<https://github.com/CityPulse/Decision-Support-and-Contextual-Filtering>).

2.6 Integration & Evaluation: Smart City Applications

Cities collect large amounts of valuable data and sensor measurements. However, a large volume of this data remains in silos and is usually used for single applications and is not often shared across different departments. The raw data streams are also not informative and hinder discovery and automated information extraction for smart city services. To overcome this semantic gap, interoperability and cross department/application sharing gaps, CityPulse annotates the data streams and employs a knowledge-based approach to interpret the data streams. The challenge is to analyse large amount of data in real-time. The CityPulse architecture splits into a very efficient large-scale data analysis pipeline and a real-time intelligence, which is fed by domain knowledge to control the data analysis in an iterative and dynamic way. Real world data is often unreliable and incomplete. CityPulse continuously monitors the quality of the information and utilises domain adaptable spatial and temporal models to improve the reliability. This allows also the detection faulty sensors and malicious sources.

Overall, the evaluation results and demos showed that even though many cities have open data portals and/or offer mechanisms to share data, the data is not used up to their its full potential unless it is provided with the right set of metadata and with suitable interfaces. The annotation models and tools provided an effective solution to address this challenge. The CityPulse applications also showed that in some cases there are gaps in the way the data is collected and shared. For the latter, CityPulse used the data stream annotation and stream processing components in the CityPulse to transform the raw data into richer information sets that could be used directly by end-user applications. The evaluations also provided interesting insights into the city ecosystems. For example, in Aarhus and by using the quality explorer, we noticed that the quality of data in some parts of the city decreases during the day. By further investigation the city of Aarhus team noticed that some of the sensors have re-chargeable batteries that use the power on lampposts. However, during the Scandinavian summer time lampposts have power only for a short period of time during the day and this did not provide enough time for the batteries to re-charge fully. Information such as this is very helpful in designing and extending the city information services. The CityPulse applications and demos also are designed for end-users and were presented and evaluated based on different feedback and insights that were collected from the users and city authorities.

2.6.1 Evaluation Scenarios

The project team performed workshops in all the three cities (i.e. Aarhus, Brasov and Stockholm) evaluating the application scenarios and prototypes developed using components from the CityPulse framework. The application scenarios and prototypes were shown to different stakeholders, city partners or citizens during a set of events using two different methodologies, one for the city of Aarhus and Brasov, and another one for the city of Stockholm. The following applications have been evaluated:

In the city of Aarhus:

- **Travel Planner** – an Android application for citizens, that can be used for obtaining user centric travel and parking recommendations;
- **3D Map** –provides a 3D visualisation of cities and different geo-located events provided by the platform;
- **Social Media Analyser** – the users can visualise events extracted from social media streams (Twitter);
- **City Dashboard** –supports visual analytics for the different relevant datasets registered in the platform;
- **Quality of Information (QoI) Explorer** – a web-based tool to get detailed insight about the quality of information that the deployed sensors provide.

In the city of Brasov:

- **Bus Travel Planner** – an Android application for citizens, that can be used for obtaining user centric bus travel recommendations
- **Brasov CKAN Server** – a data repository based on CKAN has been deployed on Brasov municipality server, useful for city administration and 3rd party developers. It already holds 4 data sets and along with the CityPulse open source components represents a starting point for developing new smart city applications.

In the city of Stockholm:

- **CityPulse Tourism Planner** – provides citizens relevant information of tourism-related points of interest and optimal routes according to the user's location, transportation information and users' preferences;
- **CityPulse Pick-up Planner** – provides optimal planning and routing for a fleet management company;
- **CityPulse Dynamic Bus Scheduler** – provides optimal bus transportation schedule based on user demands and real-time traffic in the city.

In order to evaluate the Stockholm applications/use cases a questionnaire was given to a number of participants attending the Kista Mobility Week event at Ericsson. The participants of this event were private users, academics, public sector, representatives from the telecom industry, as well as transportation companies. During this event, the applications were presented to the end users and their feedback was collected. The participants did not have any previous knowledge about the CityPulse project and the developed applications i.e. Dynamic Bus Scheduler, Tourism Planner and Pick-up Planner. The team provided a short presentation of the CityPulse project, the partners involved and listed the number of prototypes that have been developed during the project. Then we moved to presenting the applications, and we gave each participant a questionnaire to fill in. The participants had the choice to give short answers (no opinion, strongly disagree, disagree, neutral, agree, strongly agree), or to provide more descriptive feedback. The detailed results of the evaluations are published [Flein et al., 2016].

2.6.2 Overall Evaluation

The evaluation results have shown that applications are beneficial to citizens and also the city authorities and city planners in different ways including providing online information and insights, helping them to see events and occurrences around the city and plan their trip and activities. For example, the services such as Quality explorer and the 3D map provides an effective way for the city authorities to view and analyse the information that are gathered and analysed from the city in (near-) real-time. The applications such as journey planners and social data analysis among other application also present an effective way to collect and analyse the information and present live updates on events and occurrences in the city. The evaluations and discussion with citizens and different stakeholders has shown that accessing to actionable information generated by CityPulse components and applications enhances the user experience in the smart city frameworks. The open-source components also provide new opportunities and will inspire new developments and innovations in utilising the smart data and transforming the raw observations and measurements to information that will enhance the experience of citizens and/or will contribute to providing better planning and enhanced services in the cities.

During the workshops with users and citizens there were several questions regarding the end value of the products. What value could this potentially hold for me? and what value does the service represent for other perceived target groups? The presentation of the information and the amount of information was also a major topic. After having viewed the demos, there was a generally high level of enthusiasm


and the discussion reflected a detailed understanding of what the CityPulse Smart City apps were at this moment in time – show cases of what can be built on the smart city data we have available mostly for professional use. The participants were impressed with the quality of the backbone technology and development altogether and the talk among the participants reflected that the prototypes they were presented with were assessed to be the results of a high level of technical skill. The overall discussions were centered on the fact that all types of data were presented all at once.

As end users they were interested in being allowed to choose specific events and specific details. An unusual situation calls for an explanation, e.g. where do I park my car at a concert? The presentation of these filtered results should appear up front, with focus on visualisation and without all the unnecessary information.


The level of simplicity depends on the user group. Professionals were more back-end oriented and the students very preoccupied with the just-right amount of information. They were in doubt of whether the services in themselves is new or newsworthy, the quality of the visualisation and a more thoughtful presentation of the dataset is what matters if this should present a value other than what Google maps already does.

One of the key insights from discussions and evaluating the results with the stakeholders was gaining information and finding better knowledge in terms of the design and implementation of smart services in the cities. For example, during the evaluations and discussions we noticed that the quality of data in the centre of Aarhus decreases over time during the day (this was obtained by using and evaluating the QoI explorer). Figure 8 shows a picture of variation of quality of data in Aarhus (red/purple circles show lower quality and green circles show higher quality data). After further investigations we noticed that some of these sensors are charged by using the power from lampposts and during the summer because the lampposts are powered only for a very short period of time, the sensor boards do not charge fully and during the day they lose their power. The latter affects the overall quality of the data that is collected from those areas. Overall, the CityPulse experience and using the applications and services in the CityPulse showed a very good indication of effectiveness of user-led and co-design and planning in smart city applications and services.

The following present a few example of the direct feedback that are obtained from the users. The complete set of evaluations are presented in [Flein et al., 2016] and the CityPulse 3rd Annual Report [Barnaghi et al., 2016-2].



CITY: VIENNA
NAME: DR. JULIA GIRARDI-HOOG
PROJECT MANAGER, MA 25,
CITY OF VIENNA




CITYPULSE
 TRAFFIC PLANNER


The EU-funded project “Smarter Together” focuses on smart urban renewal in a neighbourhood southeast of the city centre. The aim of the project is to prepare neighbourhoods which were planned and built in the post-war period for the future. This includes infrastructure such as electricity, heating and lighting, as well as buildings, mobility, IT infrastructure and most importantly the “users” of the city. CityPulse has an interesting approach, as it is based on several kinds of data, including data produced by users. Especially for mobility purposes, it is interesting to integrate data produced real-time by users

in the neighbourhood. Also for future energy planning and energy saving, the possibility to aggregate and visualise energy production and consumption data is important.

For “Smarter Together” we chose a neighbourhood with a population that has a rather low level of formal education, a high percentage of foreign nationalities and low income. It will be crucial to find ways to integrate these users into projects like CityPulse. Certainly, the possibility to switch between language settings increases the usability of such projects.



CITY: UPPSALA
NAME: NDEY ISHA JOBE
PROFESSION: TEACHING ASSISTANT AT UPPSALA UNIVERSITY



CITYPULSE
 PICKUP PLANNER

The CityPulse Pick-up Planner is great and quite essential. It solves problems usually associated with friends (or a group of people) travelling to the same destination and at the same time wanting to reduce the cost of the trip. Uppsala, being a student city, would benefit from such a service, since students frequently need to travel to the airport to catch flights. Instead of posting on social media to inquire if their friends would like to make the same trip and consequently share the taxi cost, the students would

simply register for the pickup service and include their preferences in the requests submitted to the system.

Additionally, the service would be a great help to tourists arriving at new places and who want cheap and shared trips to various destinations. In my opinion, the best feature is the way the application optimises the route for drivers based on the users’ requests. I would definitely recommend it to any person.



CITY: VIENNA
NAME: PROF. DR. THOMAS EITER
PROFESSION: PROFESSOR OF COMPUTER SCIENCE /
HEAD OF THE KNOWLEDGE BASED SYSTEMS GROUP /
HEAD OF THE INSTITUTE OF INFORMATION SYSTEMS



CITYPULSE
TRAFFIC PLANNER

Travel planning and mobility are issues of growing demand and importance, and for smart city environments among the core problems that need attention. The increasing availability of real-time sensor data as well as, on the other hand, more dynamic background information opens up new possibilities for travel planning that consider semantic constraints, user preferences and needs, as well as multi-modal transport in order to find the best solutions. It has shown in the traffic planner demo how the framework developed in the project can be fruitfully used to

meet these problems. Based on the experience of our group at TU Wien in intelligent route planning, we see that CityPulse provides a powerful platform that by its architecture, the interfacing of Open Data, the use of semantic technologies, and the support of advanced information management and event detection, can serve as a promising tool to obtain the best solutions for travel planning tasks at increasing complexity levels. Combined with monitoring and optimisation from a global perspective this is thus right on target to enable smart cities of the future.



CITY: ATHENS
NAME: DIMITRIOS GEORGIPOULOS
PROFESSION: BUS DRIVER



CITYPULSE
DYNAMIC BUS SCHEDULER

Traffic is one of the major factors that affect the normal schedule of public transportation, especially in over-populated cities with huge numbers of private vehicles and lack of cycling infrastructure. As a bus driver, I usually receive complaints from unsatisfied passengers, when the bus does not reach its destination at the predefined time. So far, there is no solution to this problem since bus routes include the most popular roads, where traffic density is usually high, and bus drivers are obliged to follow the predefined routes, even if they might be aware of alternative and less time-consuming ones.

Based on my experience, using an application like the Citypulse Dynamic Bus Scheduler, which is able to monitor the levels of traffic density and make adjustments to the predefined bus routes, would prevent the bus vehicles from getting stuck in traffic and reduce the corresponding delays. Moreover, there would be another reason for passengers to be satisfied, since the application generates the timetables and estimates the number of operating vehicles according to their requests. And the more satisfied the passengers are, the less worried the bus driver is.



CITY: VIENNA
NAME: UNIV.-PROF. DR. AXEL POLLERES
PROFESSION: HEAD OF THE INSTITUTE
FOR INFORMATION BUSINESS AT WU WIEN
(VIENNA UNIVERSITY OF ECONOMICS & BUSINESS)



CITYPULSE
TRAFFIC PLANNER

With the traffic data analysis demo based on its platform, CityPulse provides an extensible tool that allows integrating locally fixed sensor streams and their data with other data such as mobile sensors, citizen sensing, etc. As a platform, the system – being based on Linked Data and Semantic Web technologies – uses and relies on open standards, which guarantees transparency and extensibility, plus the potential to provide and integrate Open

Data, which is at the moment one of our main research interests at the Institute of Information Business at WU Vienna. We see the potential of the combination of Open Data and Big Data – facilitated through platforms like CityPulse – as one of the future trends of data-driven economies and communities.

3 Project Impact, Dissemination and Exploitation of Results

During the lifetime of the project, the CityPulse team developed 19 open-source tools and software, defined and published 101 smart city scenarios, hosted and presented demonstrations, workshops, tutorials and keynote speeches at numerous high profile global ICT events, trained 6 PhD students, and published 11 journal articles and 26 conference papers. The partners have also commercialised extended versions of two of the developed solutions and the project's industrial partners will use the results of the projects in their future demonstrators and products.

Using the 101 scenarios, which are based on a scalable design for the smart city data analytics framework architecture, the technical team identified a set of challenges and necessary software components. The industrial and research teams at Ericsson, Siemens, the Alexandra Institute, the University of Applied Sciences in Osnabrück, the National University of Ireland in Galway, the University of Surrey and Wright State University developed a comprehensive set of open-source components and tools with several demonstrators that use these components. These demonstrators, enabling tools and software are open-source and all are available on the CityPulse GitHub <https://github.com/CityPulse/>. Some of these solutions, such as the 3D smart city map, have already become part of commercial tools or demonstrators offered by the project's industrial partners.

The open-source software, demonstrators, datasets, scenarios, documents, presentations, papers and reports available on the project website <http://www.ict-citypulse.eu/>, will not only encourage the development of smarter and more enhanced data-driven and analytical applications and services, but also make these applications more cost-efficient, agile and easier to test and try.

Key Figure	Number
<i>Scenarios Created</i>	101
<i>Components Created</i>	19
<i>Commercialised Products</i>	2
<i>Talks and Presentations Given</i>	45
<i>Outreaching Events</i>	30
<i>PHD's educated</i>	6
<i>Partners in CityPulse</i>	10
<i>Participating EU Countries</i>	6
<i>Higher Education and Research Partners</i>	4
<i>City Partners</i>	2
<i>SME</i>	1
<i>Industry Partners</i>	3
<i>Person Months Committed</i>	396

Table 2: CityPulse Results in Numbers

3.1 CityPulse Value Proposition in the Context of Smart Cities

The CityPulse framework provides a continuous and dynamic set of data analytics solutions for smart city data. The CityPulse components can be used by city administrations and third-party developers, enabling them to rapidly develop a variety of applications with a high degree of re-configurability and re-usability.

The framework has a modular structure and various components can be involved to achieve the goal. The components offer the following main functionalities:

- Data gathering and integration: Any existing or new available data coming from city streams is transformed and aligned to the CityPulse data/information model;

- Data summarisation: large amounts of data and numbers are compressed and synthesised to provide easy-to-understand and easy-to-process information;
- Data processing, in order to identify a broad range of relevant events happening in the city, in real time;
- Context based event filtering, which dynamically selects and filters events that are relevant for a particular user, in a particular location, at a particular time and also assess their level of criticality to indicate whether some actions are required in order to improve a specific activity or task the user might be engaged with;
- Decision support, which provides a set of predefined modules to help users getting recommendations and help for decision making, based on their preferences and constraints.

Application developers can use the CityPulse framework to easily deploy new data sources or migrate a smart application to a new city; since the framework supports the integration of heterogeneous data sources - both in terms of formats (e.g. two temperature streams coming from two providers and having different formats can be very easily integrated and converted to CityPulse standard format) and in terms of association with a city morphology (e.g. traffic or parking information can be processed equally from one city to another since they can be geographically linked to a uniform representation of the city roads grid).

A key aspect of observing the pulse of a city through the CityPulse framework is the concept of *events*. Events are detected, estimated, filtered and characterised based on their relevance for an end user (e.g. citizen) in the specific application domain, including traffic and mobility, cultural and social, health and environment. This is achieved via a generic event detection component where the developer can plug in various processing logic from a suit of event detection mechanism already developed. In addition to this, application developers can deploy custom-made event detection mechanisms via an easy to use application programming interfaces (APIs).

The framework offers increased stability (despite the instability of smart city data sources) by continuously monitoring the quality of the incoming data. If the quality of a data stream drops, automatic mechanisms are activated to find suitable input replacements, in order to guarantee that the information delivered via the services is always reliable and accurate. We refer to these mechanisms as Quality Monitoring and Fault recovery, which might include temporary estimation of sensor observations when the dynamic data is not received, or automatic selection of an alternative (and more reliable) data stream.

The CityPulse framework provides immediate and intuitive visual access to the results of its intelligent processing and manipulation of data and events. The ability to record and store historical (cleaned and summarised) data for post-processing makes it possible to analyse the events and occurrences in a city not only on the go but also at any point in time, enabling diagnosing and “post mortem” analysis of any incidents or relevant situation that might have occurred. To facilitate that, a dashboard for visualising the dynamic data of the smart cities is provided on top of the CityPulse framework. By using this dashboard, a city administrator has the possibility to visualise a holistic and summarised view of data across multiple contexts or a detailed view of data of interest, as well as to monitor the city life as it evolves and as things happens. The investigation of past city events or incidents can be done from different perspectives (by observing the correlations between various streams), because the streaming data is stored in the framework for a configurable period of time and can be retrieved for visualisation and analysis at any moment.

In order to demonstrate the richness and potential of CityPulse to have an impact on the cities and their inhabitants’ quality of life, we identified over 101 scenarios in different domains (e.g. energy, culture,

mobility and transportation, and health). Some of these were used to build prototype applications that demonstrate the functionality of the CityPulse framework.

New smart city applications can be envisioned and designed by consulting the list of 101 scenarios. In order to be accepted, the new technologies and applications envisioned by development companies have to be properly presented and demonstrated to the city administrators and citizens. Even if this sounds simple, it is very complicated because in most of the cases the development companies don't have direct access to the city data. In this situation, the historic data collected by the consortium becomes valuable, since it can be easily replayed (by properly configuring the CityPulse framework and not only). The historic data was collected with the help of the CityPulse pilot cities from a wide range of open data source. There is no personal data collected and as a result of that there is no risk of violating any privacy rules. This feature allows the applications development, test, presentation and demonstration without having direct access to the city data.

3.2 Impact

The CityPulse project ran over 3 years. The first year of the project was mainly about creating awareness, collecting the requirements, defining use-cases and building the technical infrastructure. The project team, in collaboration with citizens, the stakeholder group, and the technical team defined the 101 smart city use-case scenarios that were published and made available online. The project team also defined the key components and architecture of a smart city framework that was designed based on the Architecture Reference Model (ARM) for the IoT.

The efforts in the second year were mainly focused on implementation, integration and testing the proposed smart city framework. Working with the city partners in Aarhus and Brasov, the project team developed software components for real-time data stream publication and annotation, accesses and query interfaces, semantic reasoning and information processing mechanisms, adaptive analysis methods for pattern recognition and correlation analysis, and social media data analysis. The integration of these components in the framework is currently being undertaken as a smart city event analysis application. Showcasing the results and working with our project partners and receiving feedback at the development stage helped us to understand the priorities of the citizens and city authorities. Dependability and reliability of the proposed solutions, privacy and simplicity have been the key issues that we needed to take into account. We also learned several valuable lessons by involving a multi-disciplinary team in our development efforts.

The third year of the project focused on trials/evaluations, open-source and open framework development and third-party uptake and impact. The success of the project is ultimately measured on the effectiveness of the developed smart city framework and the demonstrators. This involved demonstrating the results to citizens and city stakeholder, disseminating the results, and showing the impact of the smart applications developed using the framework components in improving the access to information and services around the cities. The dissemination and exploitation activities also focused on engaging with wider group of stakeholders to create long-term impact.

Industry impact

At European level, most industrial players are more and more concerned in shifting from descriptive/diagnostic analytics to predictive/prescriptive analytics. Regardless of the domain (e.g. cities, energy, manufacturing, transportation etc.), this change will allow industrial players and their customers to reduce the costs, provide better products/services and increase the efficiency of their assets. The CityPulse generic framework provides a significant potential to be deployed/instantiated for a wide range of application domains under different complexity configuration in order to fulfil one or several analytics procedures as mentioned above. The capabilities of the framework can be used to obtain the various types of analytics:

- Descriptive/diagnostic analytics: data acquisition and storage, event detection, visualization component;
- Predictive/prescriptive analytics: decision support, technical adaptation, context based event filtering.

This allows industries to focus on the applications' goal, thus enabling rapid prototyping and deployment. Moreover, the 101 scenarios provide a great insight into the smart city landscape – helping industry to decide and prioritise their ideas. The reference datasets and demonstrators are also valuable, in particular to reach out for potential customers.

Some of the components designed in the project have already been extended and/or licenced to third-party industries for commercial deployment and development. The ontology validator developed by the University of Surrey team is being commercialised (given with a development license) by a third party company that is developing a product using this tool. The 3D map, developed at the Alexandra Institute, is capable of subscribing to extracted event and information feeds provided by other components in the project, it can also subscribe to data feeds via the extended federation component that is also offered as a commercial tool to be used by city authorities and planners. The current version includes live information for 3 cities in Denmark. However, as the user community for the CityPulse GitHub grows and as project partners involve themselves in development and innovation projects that rely on the CityPulse software components, the long-term impact of the project will become more significant.

City and Citizen Engagement

During the project life-time, the CityPulse application prototypes were shown to different stakeholders, city partners and citizens during a set of events/workshops in a form of face-to-face meetings or demos. The main workshops (but not only) were:

- Aarhus (September 25th 2015) having representative participants from City of Aarhus Services and Environment, City of Aarhus Health and Care and The central Denmark Region;
- Brasov (September 26th 2015) where the head of the IT office and representatives of the technical and mobility departments from the city hall of Brasov were invited;
- Brasov (August 25th 2016) where the attendance was made of 12 participants, representing the main city stakeholders: Representatives of the local public transport company; Representatives of the IT, road safety and transport departments of the city hall; local SMEs, members of the local IT cluster;
- Aarhus (September 15th 2016) having representative participants by two people from Smart City related SMEs and one person from The City of Aarhus.

The participants were impressed by the quality of the backbone technology and development altogether and the conversations among the participants reflected that the prototypes they were presented were assessed to be the results of a high level of technical skill. The CityPulse Annual Report for 1st, 2nd and 3rd Years have included several interviews, feedback and user evaluations for the CityPulse project. These reports are available online via the CityPulse Website and were distributed to wider research and development community that have been working with project consortium members.

Educational Activities

The project team and especially the academic members have used the results of the project and application scenarios in their teaching and educational material. For example, the University of Surrey offers an MSC/Undergraduate course on the Internet of Things to Electronic Engineering and Computer Science students. The results of the CityPulse project and in particular the open datasets and software components are used in the practices and are used to offer final year student projects. The project team

delivered several tutorials and guest lectures in international conferences and workshops on the smart city data analytics based on the results and experiences obtained in the CityPulse project.

The CityPulse demos have been used in talks given to high school students to encourage them to study science and familiarise them with the concepts of smart cities and the Internet of Things. With a rich set of demonstrators developed in the project and with several ready to use tools and presentations, the project team have been extensively involved in activities to engage public in the smart city activities and increase awareness and interests in this domain. These activities also included organising workshops and seminars to publicise and demonstrate the results of the project and plan for future developments.

Research and development

The research and development results of the project are described in detail in different sections of this report. However, overall, the project has resulted in several novel ideas, algorithms and methods to process and analyse the smart city data. The ideas and proposed solutions have been evaluated with real work data and have been published and presented to wider research community via high impact journal and conference publications. The research and industrial partners have adopted some of these solutions in their internal products and further developments. The publication and development results are described in the following sections.

3.3 Open Source Tools

The 19 open-source tools are available to be downloaded from the CityPulse GitHub account (see Figure 34): <https://github.com/CityPulse> the list of open-source tools contains:

- **CityPulse Journey Planner:** provides a travel-planning solution which goes beyond the state-of-the-art solutions by allowing users to provide multi-dimensional requirements and preferences such as air quality, traffic conditions and parking availability. In this way the users receive parking and route recommendations based on the current context of the city. In addition to this, Travel Planner continuously monitors the user context and events detected on the planned route. User will be prompted to opt for a detour if the real time conditions on the planned journey do not meet the user specified criteria anymore. All the CityPulse framework components are deployed on a back-end server and are accessible via a set of APIs. As a result of that the application developer has only to develop a user-friendly front-end application, which calls the framework APIs. In our case we have developed an Android application. Available at: <https://github.com/CityPulse/CityPulse-Journey-Planner>
- **Knowledge Acquisition Toolkit:** is a software toolkit that implements the state-of-the-art machine learning and data analytic methods for sensors data. The algorithms and methods implemented in KAT are used for processing and analysing the smart city data in the CityPulse project in the Data aggregation and Event processing components. KAT follows a simple processing approach. It uses a generic workflow that has been extracted by observing several different solutions for information abstraction. The existing solutions either follow the steps shown in the figure below or implement some parts of it. We identified the following main steps: Pre-processing to bring the data into shape for further processing, Dimensionality Reduction to either compress the data or reduce its feature vectors, Feature Extraction to find low-level Abstractions in local sensor data, Abstraction from raw data to higher-level Abstractions and finally semantic representations to make the abstracted data available for the end-user and/or machines that interpret the data. Available at: <https://github.com/CityPulse/Knowledge-Acquisition-Toolkit-2.0>
- **Data Quality Explorer:** is a web-based tool to get detailed insight about the quality of information the deploy sensors provide. Each sensor is depicted on a map as a dot. The colour of the dots in the Quality Explorer map indicates the quality of information, ranging from green

(high) to red (low). Additionally, bar charts show the number of visible sensors and the corresponding quality of information value.

Available at: <https://github.com/CityPulse/Data-Quality-Explorer>

- **Bus Travel Planner:** The citizens of Brasov can use the mobile application for receiving bus route recommendation based on the real-time location of the busses and the traffic incidents reported by the citizens. After the user has decided to travel via a certain recommended route, the system monitors the stream of incident events and if the travel preferences are not fulfilled any more, he/she is informed. The CityPulse components have been deployed on the Brasov server and are available via a set of APIs. For this prototype we have developed a user friendly application and we have properly configured the CityPulse components to work for the public transportation domain.

Available at: <https://github.com/CityPulse/Brasov-Bus-Route-Planner>

- **CityPulse Dynamic Bus Scheduler:** introduces a reasoning mechanism capable of evaluating travel requests and generating bus timetables with reduced average waiting time for passengers. Furthermore, the system has the potential to detect traffic flow and make adjustments to the regular path of each bus, so as to decrease the waiting time, which is a result of traffic congestion.

Available at: <https://github.com/CityPulse/CityPulse-Dynamic-Bus-Scheduler>

- **Resource Manager:** The Resource Management component in the CityPulse framework is responsible for managing all Data Wrappers. During runtime an application developer or the CityPulse framework operator can deploy new Data Wrappers to include data from new data streams.

Available at: <https://github.com/CityPulse/Resource-Manager>

- **Composite Data Quality Monitoring:** is used to evaluate correlations between individual data streams. It is used to check the plausibility of space-time congruent data sets. The main challenge in evaluating the correctness and information quality of heterogeneous data sources in smart city environments is a missing ground truth for comparing results. If there is no exactly planned infrastructure it is a complex process to determine, which sensor measurements are correct in case of contradictory information. This is realised in addition to the Atomic Monitoring and called Composite Monitoring. This approach helps to determine if outliers in sensor readings are due to defective sensors or can be explained by similar information from related sensors.

Available at: <https://github.com/CityPulse/Composite-Data-Quality-Monitoring>

- **Event Testing:** is an Android application for reporting events and a R Shiny application for web browsers to show and inspect events generated by the application and the CityPulse Event Detection.

Available at: <https://github.com/CityPulse/Event-Testing>

- **CityPulse Geospatial Data Infrastructure:** is used by a number of other CityPulse components to tackle geo-spatial tasks. For example, the Composite Monitoring uses the GDI to find nearby sensors or the Decision Support uses the GDI to get a set of routes across the city, which follows non-functional user requirements.

Available at: <https://github.com/CityPulse/CityPulse-Geospatial-Data-Infrastructure>

- **Event Reporting App:** enables the user to see on-going events reported by the CityPulse framework and other users. On a map view events created by CityPulse's Event Detection component and by other users of the application are shown. For user created events a possibility to change and/or delete these events is offered. The application is connected to the CityPulse framework via a message bus and requires an installed CityPulse framework. Within a configuration dialog the connection to the framework can be changed.

Available at: <https://github.com/CityPulse/EventReportApp>

- Stream Discovery and Integration Middleware:** is a middleware for complex event services. It is implemented to fulfil large-scale data analysis requirements in the CityPulse project and is responsible for event service discovery, composition, deployment, execution and adaptation. In the CityPulse project, data streams are annotated as event services, using the Complex Event Service Ontology. Event service discovery and composition refer to finding a data stream or a composite set of data streams to address the user-defined event request. Event service deployment refers to transforming the event service composition results into RDF Stream Processing (RSP) queries and registering these queries to relevant RSP engines, e.g., CQELS and C-SPARQL. Event service execution refers to establishing the input and output connections for the RSP engines to allow them consuming real-time data from lower-level data providers and delivering query results to upper-level data consumers. Event services adaptation ensures the quality constraints over event services are ensured at run-time, by detecting quality changes and make adjustments automatically.
Available at: <https://github.com/CityPulse/Stream-Discovery-and-Integration-Middleware>
- Common Libraries:** includes java classes, which represent the models of input and output for the Decision Support, Contextual Filtering, and Data Federation components in the CityPulse project.
Available at: <https://github.com/CityPulse/Common-Libraries>
- Social Media Analyser:** is composed of a dedicated Data wrapper in php (DataCollection_Phirehose-master) unit which connects to the Twitter stream API and Google translate API to simultaneously collect the data under the form of tweets and to automatically detect the source language and translate the tweets to English to facilitate the data processing step; and a data processing unit which is composed of three sub-components: a Conditional Random Field Name Entity Recognition (see below), a deep learning Convolutional Neural Network for Part of Speech tagging (see below), and a multi-view event extraction which combines the information extracted from the previous sub-components. Given a tweet and its translation, the processing unit assigns it to one of the event classes from the pre-defined class set: {Transportation and Traffic, Weather, Cultural Event, Social Event, Sport and Activity, Health and Safety, Crime and Law, Food and Drink}.
Available at: <https://github.com/CityPulse/Social-Media-Analyser>
- CityPulse 3DMap:** uses data from CityPulse components and is developed with the core goal to provide a 3D visualisation and experience to the users. By using it the users can “fly” around this 3D model of a city and visualise the effect of real-time data in the model.
Available at: <https://github.com/CityPulse/CityPulse-3D-Map>
- Fault Recovery:** ensures the continuous and proper operation of the CityPulse enabled application by generating estimated values for the data stream when the quality drops or it has temporally missing observations. When the quality of the data stream is low for a longer time period, an alternative data source has to be selected. The selection can be performed automatically by the Technical adaptation component. In other words, the technical adaptation process does not have to be triggered if the QoI of a stream is low only for a short period of time because the Fault recovery component provides estimated values. The Fault recovery component is integrated into the data wrapper. The fault recovery mechanism is triggered to generate an estimated value when the atomic monitoring component has determined that the current observation is invalid or missing.
Available at: <https://github.com/CityPulse/Fault-Recovery>
- IoT-Framework:** is a simple Restful API that provides semantically annotated data points from the IoT-Framework engine. The main requirement for using this system is an elastic search backend such as the one used by the IoT-Framework engine.
Available at: <https://github.com/CityPulse/IoT-Framework>

- Decision Support and Contextual Filtering:** enables reactive decision support functionalities to be easily deployed, providing the most suitable answers at the right time by utilising contextual information available as background knowledge, user patterns and preferences from the application as well as real-time events. The reasoning capabilities needed to support users in making better decisions require handling incomplete, diverse input, as well as constraints and preferences in the deduction process. This expressivity in the Decision Support component is achieved by using a declarative non-monotonic logic reasoning approach based on Answer Set Programming. The Decision Support component produces a set of answers to the reasoning request that satisfy all user's requirements and preferences in the best possible way. These solutions are computed by applying sets of rules deployed as scenario-driven decision support modules.
 Available at: <https://github.com/EricssonResearch/iot-framework-engine>
- CityPulse Pick-up Planner:** provides a travel service for users located around Stockholm. Users specify pickup location, destination, arrival time constraints and preferences in travel requests, from which the system devises a pickup path to be used by vehicle(s) in delivering users to their intended destinations.
 Available at: <https://github.com/CityPulse/CityPulse-Pick-up-Planner>
- Event Detector:** provides the generic tools for processing the annotated as well as aggregated data streams to obtain events occurring into the city. This component is highly flexible in deploying new event detection mechanisms, since different smart city applications require different events to be detected from the same data sources.
 Available at: <https://github.com/CityPulse/Event-Detector>
- CityPulse Tourism Planner:** combines sources of data related to events and points of interest (Pols) in the city of Stockholm, and generates a schedule to explore the Pols that the users select. The schedule is created based on the opening times of each Pol as well as the user's budget, travel period, and type of transport.
 Available at: <https://github.com/CityPulse/CityPulse-Tourism-Planner>
- CityPulse City Dashboard:** provides immediate and intuitive visual access to the results of its intelligent processing and manipulation of data and events. The ability to record and store historical (cleaned and summarised) data for post-processing makes it possible to analyse the status of the city not only on the go but also at any point in time, enabling diagnosing and "post mortem" analysis of any incidents or relevant situation that might have occurred. To facilitate that, a dashboard for visualising the dynamic data of the smart cities is provided on top of the CityPulse framework. Based on this dashboard, the user has the possibility to visualise a holistic and summarised view of data across multiple contexts or a detailed view of data of interest, as well as to monitor the city life as it evolves and as things happens. The investigation of past city events or incidents can be conducted from different perspectives, e.g. by observing the correlations between various streams, since the streaming data is stored in the framework for a period of time which can be configured, and it can be retrieved for visualisation and analysis at any time.
 Available at: <https://github.com/CityPulse/CityPulse-City-Dashboard>
- SAOPY:** is a sensor annotation library that embodies well-known ontologies in the domain of sensor networks that are used in the CityPulse project. It enables to prevent common syntax errors (e.g. undefined properties and classes, poorly formed namespaces, problematic prefixes, literal syntax) in RDF documents during the annotation process.
 Available at: <https://github.com/CityPulse/SAOPY>
- Stream Processing Benchmark:** is a java-based benchmarking toolset for RSP engines. It provides a configurable benchmark for RSP engines using CityPulse datasets (<http://deri-srvgal29.nuig.ie:8080/CityBench/>)
 Available at: <https://github.com/CityPulse/Stream-Processing-Benchmark>

- **Data-Stream-Generator:** is a tool for generating data streams. A data stream is a sequence of measurements from a sensor or other source, over time (for example temperature data, air pressure on a car's tire, level of fuel on an aircraft, etc.). It is meant to be used as a tool for testing the performance of systems designed to gather and process large amounts of data from different sources (e.g. IoT middleware). Additionally it can be used to complement existing data streams, such as the ones in the CityPulse Dataset Page
Available at: <https://github.com/CityPulse/Data-Stream-Generator>

Figure 34 shows a screenshot of the CityPulse GitHub repository (<https://github.com/CityPulse/>).

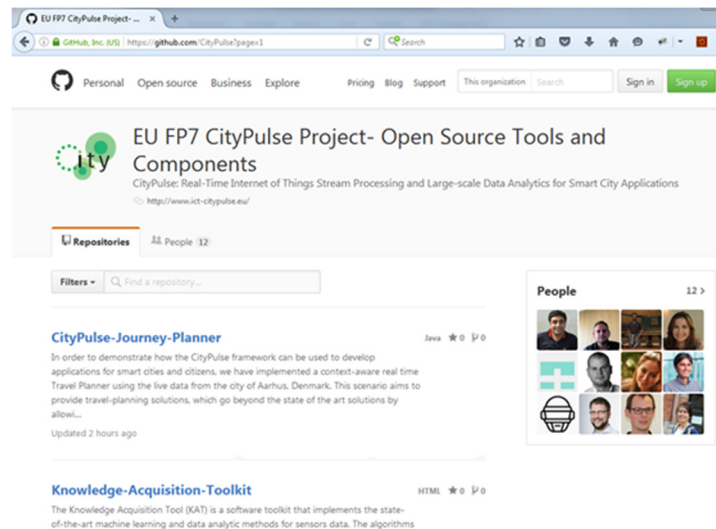


Figure 34: CityPulse GitHub

For each component, a 3rd party application developer can find a short description and a user manual in the components corresponding GitHub section. A more detailed description including the components APIs can be found in the deliverable 5.3 [Puiu et al., 2016-2].

3.4 Standardisation

During the project lifetime, the consortium members have contributed to four main standardisation bodies as follows:

- W3C Spatial Data on the Web working Group;
- W3C RDF Stream Processing Working Group;
- ITU-T Study Group 20 – Internet of Things and its applications including smart cities and communities (SC&C);
- Integration of Smart Objects: IETF CoRE, OMA LWM2M, IPSO Alliance

The CityPulse use-case scenarios are included and discussed in the requirement analysis report for W3C Spatial Data on the Web working Group. In particular, two use-cases are chosen to provide more details of our approach in the CityPulse project.

The Spatial Data on the Web Working Group <http://www.w3.org/2015/spatial/wiki/> at W3C works on describing the relevant standards and defining best practices to publish and use/reuse spatial data on the Web. The group works on determining how spatial information can be integrated with other data on the Web. It also works on describing how machines and people can discover that different facts in different datasets relate to the same place and discusses and evaluates existing methods and tools and then creates a set of best practices for their use. The CityPulse use-case scenarios are included and

discussed in the requirement analysis report for this working group. In particular, two use-cases are chosen to provide more details of our approach in the CityPulse project. The use cases are described at:

https://www.w3.org/2015/spatial/wiki/Working_Use_Cases#.22101.22_Smart_City_Use-cases

CityPulse consortium members played an important role in defining the goals and topics of the W3C RDF Stream Processing Working Group, particularly, i) RDF Stream Model, ii) RDF Query Semantics, iii) Complex Event Processing (CEP) Support for RSP, and iv) RSP Benchmarking.

CityPulse's work on the online scenarios tool it is an important reference for relevant IoT scenarios applied to smart cities by ITU-T Study Group 20.

Semantic interoperability of IoT device data and services sustains the integration of data into the Smart City Framework, and it is one of the emerging key venues for IPSO Alliance.

3.5 Dissemination of Results

As previously mentioned, throughout the project CityPulse partners successfully published their work in numerous high profile journals and conferences. In total, the partners had 11 journal and 26 conference papers accepted for publication. A list of all publications can be found in section 4 "Use and Dissemination of Foreground."

One of the most important journal papers published is "**CityPulse: Large Scale Data Analytics Framework for Smart Cities.**" The CityPulse framework developers have joined the forces and created this journal paper with the scope of providing an overview about what the CityPulse framework means and how can be used. The paper was published in IEEE Access Volume 4 (1.27 impact factor) in the dedicated section for Smart Cities. The paper has open access and at September 28th 2016 already had 846 full text views.

The project was consistently visible at numerous global ICT events throughout the project's duration, with partners delivering an extensive range of demonstrations, workshops, tutorials and talks all of which can be found in section 4.

3.6 Public website

The CityPulse website provides general information about the project, promote workshops, events and publications related to the project and continuously showcase the progress of CityPulse. The page is split into eight subsections:

- **Events:** This page shows all upcoming and past events organised by the CityPulse consortium in relation to the project.
- **Workplan:** Here the project is described in more detail. First, background information about smart cities is presented. The goals and objectives of the project are stated and a high-level view of the proposed framework is given.
- **Dissemination:** On this page all publications about are listed. They include research carried out under the funding of the project and published in various journals, conferences and workshops as well as promotions of the project through newspaper articles.
- **Tools and Data Sets:** In the project, a smart city dataset has been gathered and published online; this provides a unique set of smart city data for other researchers and developers in this field to test and evaluate their ideas and compare them with the CityPulse results. The dataset is presented in an external page linked to the main website. The size of the datasets is around 200 GBs. Furthermore, all the open-source tools and software that have been developed in the

CityPulse project are presented and their web resources are linked. All the components and tools are open source and can be found on the CityPulse GitHub page.

- **Use Cases:** Here, the requirements and metrics used to evaluate the smart city scenarios are presented. A link to a detailed description of all 101 use cases can be found here.
- **Consortium:** All participating project partners are introduced here
- **Links:** A collection of related Internet resources.
- **Internal:** Links to the internal wiki of the project, where the components and the progress is documented in detail. The accessibility is only granted for people working on the project.
- **Contact:** The contact details of the project coordinator, project administrator, city stakeholder group coordinator and privacy officer are given here.

3.7 Beyond CityPulse

Partners will continue to build on and exploit the results of the project to ensure a lasting legacy for CityPulse. Some examples of the current and ongoing exploitation of the project's results are:

- Aspern Living Labs: Will potentially use some of the CityPulse components in a follow up project. More precisely, the SAOPY library, together with the Stream annotation ontology is being considered for annotating the sensory data generated by the smart buildings and the smart grid, and both the Atomic and Composite Data Quality Monitoring components can be useful to detect faulty sensors.
- Siemens City Intelligence Platform (CIP): There are currently plans to extend the CIP with more fine-grained, personalised decision support features, based on particular situations. For that, the Decision Support and Contextual Filtering component from CityPulse is being considered.
- Future exploitation involving different cities: SAGO and Siemens Germany have joint projects involving cities such as Munich, Paris, Bilbao, Manchester, and Bratislava in the context of city sustainability and resilience, where CIP is being used.
- Further development of a Composite Data Stream Evaluation Toolbox based on evaluated models of the Composite Monitoring and a Testbed for Automated Monitoring and Testing in Smart Cities. Both tools will increase the expertise gained by PhD students and will also be available in combination with the CityPulse framework for students of the bachelor course Software Engineering Project at the University of Applied Science Osnabrück.
- The collection and analysis of datasets from the city of Aarhus will be continued to optimise the Information Quality Framework. The framework will be utilised in further research project of the IoT domain. In addition, the knowledge gained will be exchanged with the municipal works of the city of Osnabrück especially in the sector of Smart Grids.
- The KAT toolkit continues to be extended and further developed by the user community and also by the original development team. Kat is currently being extended to provide data analytics as a service and to also include extended and new set of algorithms.
- The City partners plan to run workshops and events to encourage the use and extension of framework components in new application and services provided by start-ups and innovators.
- Smart City engagement is a strategic growth for Ericsson.
- CityPulse results in terms of use cases and applications as well as architecture and the smart city framework provides Ericsson with insights on technologies that form input to portfolio planning and customer engagement.

3.8 The CityPulse Consortium



The University of Surrey
www.surrey.ac.uk/ics

The Alexandra Institute
www.alexandra.dk/uk

The University of Applied Sciences Osnabrück
www.uni-osnabrueck.de/en/home

National University of Ireland
www.insight-centre.org

Wright State University
www.knoesis.org

Ericsson
www.ericsson.com

SIEMENS Romania
www.siemens.com

City of Brasov
www.metropolabrasov.ro

City of Aarhus
www.dokk1.dk

SIEMENS Austria
www.siemens.com

4 Use and Dissemination of Foreground

4.1 Publications

4.1.1 Journal

#	Author(s)	Contribution Title	Journal/Book Name	Publication Date	Link
1	Pramod Anantharam, Payam Barnaghi, K. Thirunarayan, Amith Sheth	Extracting city events from social streams	ACM Transactions on Intelligent Systems & Technology	January 2015	http://tist.acm.org/2015vol6no4.html
2	F. Ganz, D. Puschmann, P. Barnaghi, F. Carrez	A Practical Evaluation of Information Processing and Abstraction Techniques for the Internet of Things	IEEE Internet of Things Journal	March 2015	http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7055837&filter%3DAND%28p_IS_Number%3A7173095%29
3	Payam Barnaghi, Amit Sheth, Vivek Singh, Manfred Hauswirth	Physical-Cyber-Social Computing: Looking Back, Looking Forward	Guest Editors Introduction, IEEE Internet Computing, May/June 2015	June 2015	http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7111889
4	P. Barnaghi, M. Bermudez-Edo, R. Toenjes	Challenges for Quality of Data in Smart Cities	ACM Journal of Data and Information Quality, challenge paper	July 2015	http://dl.acm.org/citation.cfm?id=2747881&CFID=700717221&CFTOKEN=50374902
5	Payam Barnaghi, Martin Bauer, Abdur Rahim Biswas, Maarten Botterman, Bin Cheng, Flavio Cirillo, Markus Dillinger, Hans Graux, Seyed Amir Hoseinitabatabaie, Ernő Kovacs, Salvatore Longo, Swaroop Nunna, Alois Paulin, R. R. Venkatesha Prasad, John Soldatos1, Christoph Thuemmler and Mojca Volk	IoT Analytics: Collect, Process, Analyze, and Present Massive Amounts of Operational Data – Research and Innovation Challenges	Book chapter (Chapter 7), in "Building the Hyperconnected Society: Internet of Things Research and Innovation Value Chains, Ecosystems and Markets, edited by Ovidiu Vermesan, Peter Friess, River Publishers	2015	http://tinyurl.com/ppljmyx
6	D. Puiu, P. Barnaghi, R. Toenjes, D. Kumper, M. I. Ali, A. Mileo, J. X. Parreira, M. Fischer, S. Kolozali, N. Farajidavar, F. Gao, T. Iggena, T. L. Pham, C. S. Nechifor, D. Puschmann, and J. Fernandes	Citypulse: Large scale data analytics framework for smart cities	IEEE Access, vol. 4, pp. 1086–1108	2016	http://ieeexplore.ieee.org/document/7447851/
7	F. Gao, M.I. Ali, E. Curry and A. Mileo	QoS-aware Stream Federation and Optimisation based on Service Composition	International Journal on Semantic Web and Information Systems (IJSWIS), 2016.	2016	http://www.igi-global.com/article/qos-aware-

					stream-federation-and-optimization-based-on-service-composition/164484
8	Sefki Kolozali et al	On the Effect of Adaptive and Non-Adaptive Analysis of Time-Series Sensory Data	IEEE Internet of Things Journal	2016	http://ieeexplore.ieee.org/document/7452350/?arnumber=7452350
9	Daniel Puschmann et al	Adaptive Clustering for Dynamic IoT Data Streams	IEEE Internet of Things Journal	2016	Accepted; to appear
10	Payam Barnaghi, Amit Sheth	On Searching the Internet of Things: Requirements and Challenges	IEEE Intelligent Systems	2016	Accepted; to appear
11	Andreas Kamilaris et al.	Mobile Phone Computing and the Internet of Things: A Survey	IEEE Internet of Things Journal	2016	Accepted, to appear

4.1.2 Conference

#	Author(s)	Title	Conference	Location	Date	Link
1	F. Ganz, P. Barnaghi, F. Carrez	Multi-resolution Data Communication in Wireless Sensor Networks	IEEE World Forum on Internet of Things WF-IoT	Seoul, Korea	March 2014	http://personal.ee.surrey.ac.uk/Personal/P.Barnaghi/doc/GANZMULTI.pdf
2	T. Iggena, D. Kümper, R. Tönjes	Kontinuierliche Bewertung von Informationsqualität in Stream-basierten Smart City Architekturen	ITG-Fachbericht Mobilkommunikationstechnologien und Anwendungen	Osnabrück, Deutschland	May 2014	https://www.vde-verlag.de/proceedings-en/453611010.html
3	F. Gao, E. Curry, S. Bhiri	Complex Event Service Provision and Composition based on Event Pattern Matchmaking	8th ACM International Conference on Distributed Event-Based Systems, DEBS'14	Mumbai, India	May 2014	https://www.insight-centre.org/content/complex-event-service-provision-and-composition-based-event-pattern-matchmaking
4	R. Tönjes, P. Barnaghi, M. Ali, A. Mileo, M. Hauswirth, F. Ganz, S. Ganea, B. Kjærgaard, D. Kuemper, S. Nechifor, D. Puiu, A. Sheth, V. Tsiatsis, L. Vestergaard	Real Time IoT Stream Processing and Large-scale Data Analytics for Smart City Applications	European Conference on Networks and Communications 2014	Bologna, Italy	June 2014	http://2014.eswc-conferences.org/sites/default/files/eswc2014euprojects_submission_17.pdf
5	S. Bischof, A. Karapantelakis, A. Sheth, A. Mileo, S. Nechifor, P. Barnaghi	Semantic Modelling of Smart City Data	W3C Workshop on the Web of Things, Enablers and services for an open Web of Devices	Berlin, Germany	June 2014	http://www.w3.org/2014/02/wot/Overview.html
6	K. Vandikas, V. Tsiatsis	Performance evaluation of an IoT Platform	Eighth International Conference on Next Generation Mobile Apps, Services and Technologies, NGMAST 2014	Oxford, UK	September 2014	http://ieeexplore.ieee.org/xp/articleDetails.jsp?tp=&arnumber=6982906&queryText=%3Dvandikas

7	Septimiu Nechifor, Dan Puiu, Bogdan Tarnauca, Florin Moldoveanu	Autonomic aspects of IoT based systems. A logistics domain scheduling example	Workshop on Interoperability and Open-Source Solutions for the Internet of Things	Split, Croatia	September 2014	http://openiot.eu/?q=softcom2014workshop
8	F. Gao, M. I. Ali, A. Mileo	Semantic Discovery and Integration of Urban Data Streams	5th Workshop on Semantics for Smarter Cities	Riva del Garda, Italy	October 2014	http://ceur-ws.org/Vol-1280/paper5.pdf
9	M. I. Ali, A. Mileo	How good is your SPARQL endpoint? A QoS-Aware SPARQL Endpoint Monitoring and Data Source Selection Mechanism for Federated SPARQL Queries	13th International Conference on Ontologies, DataBases, and Applications of Semantics, ODBASE 2014	Amantea, Italy	October 2014	http://link.springer.com/chapter/10.1007%2F978-3-662-45563-0_29
10	F. Gao, E. Curry, M. I. Ali, S. Bhiri, A. Mileo	QoS-aware Complex Event Service Composition and optimization using Genetic Algorithms	12th International Conference on Service Oriented Computing, ICSOC 2014	Paris, France	November 2014	http://www.ict-citypulse.eu/page/sites/default/files/icsoc.pdf
11	A. Blázquez, V. Tsiatsis, K. Vandikas	Performance Evaluation of OpenID Connect for an IoT Information Marketplace	International Workshop: Data Analytics for Dynamic Environments (DADE 2015)	Glasgow, Scotland	May 2015	http://ieeexplore.ieee.org/document/7146004/authors
12	A. Moregård, K. Vandikas	Computations on the Edge in the Internet Of Things	The 6th International Conference on Ambient Systems, Networks and Technologies (ANT 2015)	London, United Kingdom	June 2015	https://www.diva-portal.org/smash/get/diva2:867715/FULLTEXT01.pdf
13	S. Bischof, C. Martin, A. Polleres, P. Schneider	Collecting, Integrating, Enriching and Re-publishing Open City Data as Linked Data	The 14th International Semantic Web Conference (ISWC 2015)	Bethlehem, U.S.	October 2015	http://iswc2015.semanticweb.org/
14	R. Toenjes, D. Kuemper and M. Fischer	Knowledge-Based Spatial Reasoning for IoT-Enabled Smart City Applications	2015 IEEE International Conference on Data Science and Data Intensive Systems	Sydney, NSW, Australia	December 2015	http://ieeexplore.ieee.org/document/7396583/?arnumber=7396583&tag=1
15	Pramod Anantharam, Krishnaprasad Thirunarayan, Surendra Marupudi, Amit Sheth, Tanvi Banerjee	Understanding City Traffic Dynamics Utilizing Sensor and Textual Observations	The Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)	Phoenix, Arizona, USA	February 2016	http://www.knoesis.org/library/resource.php?id=2228
16	T. Iggena, D. Kümpfer, M. Fischer, R. Tönjes.	Monitoring and Testing for Reliable Smart City Applications	21. ITG Fachtagung Mobilkommunikation, Osnabrück	Osnabrueck, Germany	May 2016	https://www.vde-verlag.de/proceedings-de/454220016.html
17	Lucian Sasu, Dan Puiu, Septimiu Nechifor	Fault recovery mechanism for smart city environments	Intelligent Engineering Systems (INES), 2016 IEEE 20th Jubilee International Conference on	Budapest HU	June 2016	http://ieeexplore.ieee.org/document/7555093/
18	Shen Gao, Daniele Dell'Aglio, Soheila Dehghanzadeh, Abraham Bernstein, Emanuele Della Valle and Alessandra Mileo	Planning Ahead: Stream-Driven Linked-Data Access under Update-Budget Constraints.	International Semantic Web Conference, 2016	Kobe, Japan	October 2016	http://iswc2016.semanticweb.org/pages/program/accepted-papers.html

19	Septimiu Nechifor, Ioana Stefan, Marten Fischer, Dan Puiu	Event Detection for Urban Dynamic Data Streams	2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW 2016)	Barcelona, Spain	December 2016	Accepted for publication
20	D. Kümper, M. Fischer, T. Iggena, R. Tönjes, E. Pulvermüller	Monitoring Data Stream Reliability in Smart City Environments	2016 IEEE 3rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
21	Daniel Puschmann et al	Marginal Distribution Clustering of Multi-variate Streaming IoT Data	2016 IEEE 3rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
22	Yasmin Fathy, et al.	A Distributed In-network Indexing Mechanism for the Internet of Things	2016 IEEE 3rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
22	Nazli Farajidavar et al.	Physical-Cyber-Social Similarity Analysis in Smart Cities	2016 IEEE 3rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
24	Andreas Kamilaris, M.I. Ali	Do "Web of Thing" Truly Follow the Web of Things?	2016 IEEE 3 rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
25	M.I. Ali et al.	WOTS2E: A Search Engine for a Semantic Web of Things	2016 IEEE 3 rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication
26	Andreas Kamilaris, Feng Gao, M.I. Ali	Agri-IoT: A Semantic Framework for Internet of Things-enabled Smart Farming Applications	2016 IEEE 3 rd World Forum on Internet of Things WF-IoT	Reston, VA, USA	December 2016	Accepted for Publication

4.2 Dissemination Activities

#	Type of Activity	Main Contributor(s)	Description / Title	Location	Date	Link
1	Workshop	Siemens, Brasov Metropolitan Agency, The Environment Agency of Brasov County	Scenario Workshop	Brasov, Romania	September 2013	
2	Workshop	Siemens, Brasov Metropolitan Agency, Brasov City Hall	Scenario Workshop	Brasov, Romania	October 2013	
3	Workshop	CSG and project partners	Scenario Workshop	Stockholm, Sweden	October 2013	
4	Workshop	Alexandra Institute and Aarhus Municipality	Scenario Workshop	Aarhus, Denmark	November 2013	
5	Presentation	Sorin Ganea/Braşov Metropolitan Agency	CityPulse: Real-Time IoT Stream Processing and Large-scale Data Analytics for Smart City Applications	Brasov, Romania	November 2013	http://www.metropolabrasov.ro/proiecte.php
6	Workshop	Alexandra Institute and Aarhus Municipality	Scenario Workshop	Aarhus, Denmark	December 2013	
7	Workshop	CSG, City of Vienna – MA 18, ASCR Vienna, Siemens	Scenario Workshop	Vienna, Austria	December 2013	

		AG Austria, INSIGHT Galway, University of Surrey				
8	Presentation	Payam Barnaghi	IoT as a proponent of new Business Models and Social Engagement in Smart Cities	FIA, Athens, Greece	March 2014	http://www.fi-athens.eu/program/workshops/iot-proponent-new-business-models-and-social-engagement-smart-cities
9	Presentation	Payam Barnaghi	Internet of Things: The story so far	Invited talk at Conference on Rolling Plan on ICT Standardisation, Brussels, Belgium	March 2014	http://www.slideshare.net/citypulse/internet-of-things-the-story-so-far-32660875
10	Workshop	Siemens, Brasov Metropolitan Agency, Brasov City Hall	Scenario Workshop	Brasov, Romania	March 2014	
11	Workshop	Alexandra Institute, CISCO, Aarhus University and Aarhus Municipality	Scenario Workshop	Aarhus, Denmark	April 2014	
12	Presentation	Wright State University	Transforming Big Data into Smart Data: Deriving Value via Harnessing Volume, Variety, and Velocity Using Semantic Techniques and Technologies	Keynote speaker at the 30th IEEE International Conference on Data Engineering, Chicago, Illinois	April 2014	http://bit.ly/1Mw3dAj
13	Presentation	Michelle Bak Mikkelsen	Aarhus Data Drinks #6 Open Data and Hacktivism.	Aarhus, Denmark	April 2014	https://www.facebook.com/events/746683755344884/
14	Presentation	Payam Barnaghi	Public Internet of Things	ISSNIP Symposium, Singapore	April 2014	http://publiciot.org/
15	Presentation	P. Barnaghi, R. Tonjes, J. Holler, M. Hauswirth, A. Sheth, P. Anantharam	CityPulse: Real-Time IoT Stream Processing and Large-scale Data Analytics for Smart City Applications	EU Project Networking Session at ESWC 2014	May 2014	http://www.ict-citypulse.eu/doc/CityPulseExtendedAbstract_ESWC_EU.pdf
16	Presentation	Michelle Bak Mikkelsen	Aarhus Data Drinks #7 Open Data and entrepreneurship	Aarhus, Denmark	June 2014	https://www.facebook.com/events/690577384336395/
17	Presentation	Payam Barnaghi	Discovering Things and Things' data/services	IoT Week, London, UK	June 2014	http://www.slideshare.net/PayamBarnaghi/discovery-iot-week2014
18	Presentation	Payam Barnaghi	Working with Real World Data	IoT Week, London, UK	June 2014	http://www.slideshare.net/PayamBarnaghi/urb-grade-dataiotweek2014
19	Presentation	Wright State University	Smart Data for you and me: Personalized and Actionable Physical Cyber Social Big Data	Keynote speaker at the 2014 World Congress in Computer	July 2014	http://www.worldacademyofscience.org/worldcomp14/ws

				Science, Computer Engineering, and Applied Computing		
20	Presentation	Payam Barnaghi	Large-scale Data Analytics for Physical-Cyber-Social Streams	Invited talk at the 4th International Joint Workshop on Cyber-Physical Cloud Computing, Osaka, Japan	August 2014	http://www.slideshare.net/PayamBarnaghi/data-analytics-for-smart-cities-39701828
21	Presentation	Michelle Bak Mikkelsen	Aarhus Data Drinks #7 Open Data and IT in healthcare	Aarhus, Denmark	September 2014	https://www.facebook.com/events/1454879654784177/
22	Media	P. Barnaghi, A. Sheth	The Internet of Things: The Story So Far	IEEE eNewsLetter	September 2014	http://iot.ieee.org/newsletter/september-2014/the-internet-of-things-the-story-so-far.html
23	Media	Braşov Metropolitan Agency	"City Pulse partners' meeting in Brasov" article in Bună Ziua Braşov	Newspaper Article	September 2014	http://www.bzb.ro/stire/intalnirea-partenerilor-proiectului-city-pulse-la-brasov-a77755
24	Media	Braşov Metropolitan Agency	"City Pulse partners' meeting in Brasov" article in Transilvania Express	Newspaper Article	September 2014	http://www.mytex.ro/agenda-judeteana-brasov/intalnirea-partenerilor-proiectului-city-pulse-la-brasov_407348.php
25	Presentation	Payam Barnaghi	CityPulse: Large-scale data analytics for smart cities	IEEE IoT 2014, Taipei, Taiwan	September 2014	http://www.slideshare.net/PayamBarnaghi/data-analytics-for-smart-cities
26	Demonstration	Lasse Steenbock Vestergaard	CityPulse Scenario 1 implementations and list of 101 scenarios	IOT360 – the gateway to Innovation. Rome, Italy	October 2014	http://iot-360.eu/2014/about/
27	Presentation	Wright State University	Course on Smart Cities – "Role of Big Data in Smart City Applications"	Centre for Environmental Planning and Technology (CEPT) University	October 2014	http://bit.ly/1h7XQey
28	Presentation	Wright State University	Smart Data - How you and I will exploit Big Data for personalized digital health and many other activities	Keynote speaker at the IEEE Big Data 2014. Bethesda, MD	October 2014	
29	Presentation	Payam Barnaghi	Dynamic Semantics for Dynamic IoT Environments	Keynote speech at The 7th International Workshop on	October 2014	http://www.slideshare.net/PayamBarnaghi/dynamic-semantics-for-semantics-

				Semantic Sensor Networks (SSN2014)		for-dynamic-iot-environments
30	Workshop	Co-hosted by CityPulse	5th Workshop on Semantics for Smarter Cities	The 13th International Semantic Web Conference (ISWC2014), Riva del Garda, Italy,	October 2014	http://blog.soton.ac.uk/s4sc/
31	Media	Braşov Metropolitan Agency	"The citizens of Brasov can evaluate the Smart Cities scenarios" article in Monitorul Express	Newspaper Article	January 2015	
32	Media	Braşov Metropolitan Agency	"The citizens of Brasov can evaluate the Smart Cities scenarios" article in Bună Ziua Braşov	Newspaper Article	January 2015	
33	Media	Braşov Metropolitan Agency	Promoting the public evaluation of the 101 scenarios "The citizens of Brasov can evaluate the Smart Cities scenarios"	Braşov Metropolitan Agency Website	January 2015	http://metropolabrasov.ro/
34	Media	Braşov Metropolitan Agency	Promoting the public evaluation of the 101 scenarios "The citizens of Brasov can evaluate the Smart Cities scenarios"	Braşov Metropolitan Agency facebook page	January 2015	https://www.facebook.com/primariabrasov
35	Presentation	Payam Barnaghi	Dynamic Semantics for the Internet of Things	Invited talk at the Ontology Summit 2015, Session - Ontology Integration in the Internet of Things	February 2015	http://www.slideshare.net/PayamBarnaghi/dynamic-semantics-for-the-internet-of-things
36	Media	Braşov Metropolitan Agency	"The citizens of Brasov are invited to evaluate the smart city type scenarios developed in the Citypulse project" article in Transilvania Express	Newspaper Article	February 2015	http://www.mytex.ro/component/content/article/170-agenda-locala/zona-metropolitana/292198-brasovenii-pot-participa-direct-la-evaluarea-scenariilor-de-tip-smart-city-elaborate-in-proiectul-city-pulse.html
37	Media	Payam Barnaghi et al.	IET Sector Technical Briefing – "Digital Technology Adoption in the Smart Built Environment"	The Institution of Engineering and Technology (IET)	March 2015	http://www.theiet.org/sector/s/built-environment/resources/digital-technology.cfm
38	Workshop	Sorin Ganea / Brasov Metropolitan Agency	How to start your start up	Braşov, Romania	May 2015	
39	Presentation	Payam Barnaghi	Physical-Cyber-Social Data Analytics & Smart City Applications	5th Annual International Cyber-Physical Cloud Computing Workshop, Washington DC, USA	April 2015	http://www.slideshare.net/PayamBarnaghi/dynamic-semantics-for-the-internet-of-things
40	Presentation	Payam Barnaghi	Data Analytics for Smart Cities: Looking Back, Looking Forward	Plenary talk at Smart Cities Convention, London	April 2015	http://www.slideshare.net/PayamBarnaghi/data-analytics-for-smart-cities-

						looking-back-looking-forward Video: https://youtu.be/5DY5pUeb1JU?t=2445
41	Presentation	Septimiu Nechifor	IoT Research Project: CityPulse	IoTWeek 2015. Lisbon, Portugal	June 2015	http://iot-week.eu/
42	Presentation	Payam Barnaghi	Dynamic Data Analytics for the Internet of Things: Challenges and Opportunities	IoT Large-Scale Analytics Workshop IoTWeek 2015. Lisbon, Portugal	June 2015	http://www.slideshare.net/PayamBarnaghi/dynamic-data-analytics-for-the-internet-of-things-challenges-and-opportunities
43	Demonstration	Dan Puiu	Project booth - CityPulse framework, Scenario 1 and list of 101 scenarios	SIDO – The connected business. Lyon, France	June 2015	http://www.sido-event.com/en/event-connected-objects-lyon-france.html
44	Workshop	Sorin Ganea / Brasov Metropolitan Agency	Alt Start	Brasov, Romania	June 2015	
45	Demonstration	Dan Puiu, Ali Intizar, Thorben Iggena	Project booth - CityPulse framework, Scenario 1 and 2	IoTWeek 2015. Lisbon Portugal	June 2015	http://iot-week.eu/
46	Presentation	Payam Barnaghi	The Future of the Internet	Keynote speaker at The Grand Challenge 2015, University of Exeter	June 2015	http://www.slideshare.net/PayamBarnaghi/the-future-of-the-internet-48910037
47	Media	Siemens AG Austria	The Digital Pulse of the City	Hitech - the technology blog of Siemens Austria	June 2015	http://www.hitech.at/2015/06/24/der-digitale-puls-der-stadt/
48	Presentation	Payam Barnaghi	CityPulse: Large-scale data analysis for smart city applications	Invited talk at the IoT & Society Workshop, IoT Week. Lisbon	June 2015	http://www.slideshare.net/PayamBarnaghi/citypulse-largescale-data-analysis-for-smart-city-applications
49	Presentation	Payam Barnaghi	CityPulse: Large-scale data analysis for smart city applications	Invited talk at the EU-US Smart Cities Workshop, IoT Week Lisbon	June 2015	http://www.slideshare.net/PayamBarnaghi/citypulse-largescale-data-analysis-for-smart-city-applications-51382495

50	Presentation	Payam Barnaghi	The Future of the Internet	The Grand Challenge 2015, University of Exeter,	June 2015	http://www.slideshare.net/PayamBarnaghi/the-future-of-the-internet-48910037
51	Presentation	Payam Barnaghi, Ali Intizar, Şefki Kolozali, Alessandra Mileo	Semantics and Data Analytics for Smart City Applications	Tutorial at 12th Extended Semantic Web Conference 2015 (ESWC)	June 2015	http://ict-citypulse.eu/page/tutorial/smartcity-tutorial-eswc2015.htm
52	Media	Braşov Metropolitan Agency	"Demonstrator software developed for the city of Braşov in the City Pulse project" article in Bună ziua Braşov	Newspaper Article	August 2015	http://www.bzb.ro/stire/aplicatie-software-cu-rol-demonstrativ-dezvoltata-pentru-municipiul-brasov-in-cadrul-proiectului-city-pulse-a89332
53	Presentation	Payam Barnaghi	Semantic Technologies for the Internet of Things: Challenges and Opportunities	MyIoT Week Malaysia 2015, MIMOS Berhad, Kuala Lumpur, Malaysia	August 2015	http://www.slideshare.net/PayamBarnaghi/semantic-technologies-for-the-internet-of-things-challenges-and-opportunities
54	Presentation	Payam Barnaghi	Opportunities and Challenges of Large-scale IoT Data Analytics	ASEAN IoT Innovation Forum, Kuala Lumpur, Malaysia,	August 2015	http://www.slideshare.net/PayamBarnaghi/opportunities-and-challenges-of-largescale-iot-data-analytics
55	Presentation	Payam Barnaghi	Internet of Things and Large-scale Data Analytics ,	The IET Surrey Network, Guildford, UK	September 2015	http://www.slideshare.net/PayamBarnaghi/internet-of-things-and-largescale-data-analytics
56	Presentation	Amit Sheth	Semantics-empowered Smart City Applications: Today and Tomorrow	Keynote speaker at The 6th Workshop on Semantics for Smarter Cities (S4SC 2015)	October 2015	http://www.slideshare.net/apsheth/semanticsempowered-smart-city-applications-today-and-tomorrow?utm_source=slideshow&utm_medium=ssemail&utm_campaign=upload_digest
57	Presentation	Payam Barnaghi	Smart Cities and Data Analytics: Challenges and Opportunities	Workshop on Smart City: Applications and Services, Budva, Montenegro	October 2015	http://www.slideshare.net/PayamBarnaghi/smart-cities-and-data-analytics-challenges-and-opportunities
58	Presentation	Michelle Bach Lindstrøm	CityPulse: Real-Time IoT Stream Processing and Large-scale Data Analytics for Smart City Applications	Aarhus Mini Market Faire. Aarhus, Denmark	October 2015	http://makerfaireaarhus.dk/146-2/

59	Presentation	Ali Intizar	Stream Reasoning: Reasoning Upon Rapidly Changing Information (SR4LD)	Tutorial at ISWC 2015, Pennsylvania, US	October 2015	http://streamreasoning.org/events/sr4ld2015
60	Presentation	Ali Intizar	Federated Query Processing Over Linked Data	Tutorial at ISWC 2015, Pennsylvania, US	October 2015	https://sites.google.com/site/qfedld/
61	Presentation	K. Taylor	Smart City and open standards	Keynote talk at OGC's Smart Cities Conference	November 2015	http://www.ict-citypulse.eu/page/sites/default/files/smart_cities_and_open_standards_1.pdf
62	Presentation	Payam Barnaghi	Semantics and Data Analytics for Smart City Applications	Tutorial at Extended Semantic Web Conference, ESWC 2015	2015	http://ict-citypulse.eu/page/tutorial/smartcity-tutorial-eswc2015.htm
63	Presentation	Daniel Kuemper	CityPulse-Reliable Information Processing for Smart City Applications	Tutorial at Mobile Systems symposium of GIN, Germany	2015	http://www.gin-online.de/index.php/8-nachrichten/25-gin-forum-mobile-systeme
64	Presentation	Dan Puiu	Large-scale Data Analytics for Smart City and Industrial Applications	IT Days. Cluj, Romania	2015	http://www.itdays.ro/
65	Presentation	Nazli Farajidavar	social media for smart healthcare	Kent Surrey Sussex Academic Health Science Network Expo and Awards, UK	January 2016	http://www.eventbrite.co.uk/e/kss-ahsn-expo-awards-2016-registration-18570960242?utm_campaign=reminder_attendees_48_hour_email&utm_medium=email&utm_source=eb_email&utm_term=eventname
66	Demonstration	Dan Puiu, Payam Barnaghi, R. Tönjes	Project booth - CityPulse Framework	Mobile World Congress Barcelona, Spain	February 2016	http://www.gsma.com/event/s/mobile-world-congress
67	Presentation	Azadeh Bararsani	Three CityPulse use-cases	Kista Mobility Week, Stockholm, Sweden	April 2016	https://www.ericsson.com/thecompany/events/kista-mobility-week-2016_957590553_c
68	Media	Konstantinos Vandikas	A recipe for personalized bus transportation	Ericsson Research blog	May 2016	https://www.ericsson.com/research-blog/smart-cities/personalized-bus-transportation/
69	Presentation	Ralf Tönjes	CityPulse: Reliable Information Processing in Smart City Frameworks	Invited talk at IoTWeek, Belgrade, Serbia	May 2016	http://www.ict-citypulse.eu/page/sites/default/files/iot-week16-toenjes.pdf

70	Demonstration	Josiane Xavier Parreira, Stefan Bischof	CityPulse demonstrators: Travel planner app, Quality Explorer and Social Media Analysis	SIEMENS Smart City Lounge. Vienna, Austria	June 2016	
71	Presentation	K. Taylor	Semantic Sensor Networks	W3C/OGC Spatial Data on the Web session. Australian National University	June 2016	http://www.ict-citypulse.eu/page/sites/default/files/acsw_february_2016.pdf
72	Media	Azadeh Bararsani	Smart mobility apps for smarter living	Ericsson Research blog	July 2016	https://www.ericsson.com/research-blog/data-knowledge/smart-mobility-apps-smarter-living/
73	Presentation	Nazli Farajidavar and Sefki Koložali	Stream processing and data analytics for smart city	Big Data and Analytics Summer School, University of Essex, UK	September 2016	https://www.essex.ac.uk/iads/documents/2016-programme.pdf
74	Presentation	Ali Intizar	2nd tutorial on RDF-Stream Processing	Tutorial at the International Semantic Web Conference 2016, Kobe, Japan	October 2016	http://streamreasoning.org/events/rsp2016
75	Presentation	Ali Intizar	Semantic Web meets Internet of Things (IoT) and Web of Things (WoT)	Tutorial at the International Semantic Web Conference 2016, Kobe, Japan	October 2016	http://sensormeasurement.appspot.com/?p=ISWC2016Tutorial

4.2.1 Exploitable Foreground

Type of Exploitable Foreground	Description of exploitable foreground	Confidential YES/NO	Foreseen embargo date	Exploitable product(s) or measure(s)	Sector(s) of application	Timetable, commercial or any other use	Patents or other IPR exploitation (licences)	Owner & Other Beneficiary(s) involved
Software	Open source software of a 3D map and data visualizations for Smart Cities	No	N/A	Custom 3D map installations; integration of map in existing digital services	Smart City	Already a commercial product	Open Source	AI
Scenario Tool	Use Case or Scenario based Ideation Tool	No	N/A	101 Scenarios in a Digital form	Smart City	Links to other projects (OrganCity) and Organisations (IEEE); used in consultancy work	Open Access	AI
QoI-Monitoring and Explorer	Open Source Tool to monitor quality of spatiotemporal distributed data	No	N/A	Support for Tool adaptations. Integration into existing IoT / Smart City platforms	IoT and Smart City	Presentation of extended tool in 2017 (Composite Monitoring approach)	Open Source	UASO
Software	Stream Discovery and Integration Middleware	No	N/A	Data Federation Component for IoT/Smart City Platforms	IoT, Smart Cities, and Smart Farming	IoT data portal for Ireland	Open Source	NUIG
Software	Decision Support and Contextual Filtering	No	N/A		IoT and Smart Cities		Open Source	NUIG
Software	Brasov Bus Planner	No	N/A	Planning application	Smart city / smart transportation	Deployment for city of Brasov already done	Open Source	SIEMENS RO and AT
Software	Event detection	No	N/A	Event extraction tool	Applicable in any domain	Already integrated in SIEMENS BAM portfolio	Open Source	SIEMENS RO

Software	Dashboard	No	N/A	Data visualization tool	Applicable in any domain	Already integrated in SIEMENS BAM portfolio	Open Source	SIEMENS RO
General advancement of knowledge	Knowledge about real-time data analytics and smart cities	No	N/A	The knowledge acquired during the CityPulse framework development and test has the potential to improve the next generation of SIEMENS data analytics products	Applicable in any domain	N/A	N/A	SIEMENS RO
Software	CityPulse Tourism Scheduler	No	N/A	Planning application	Smart city / Smart transportation	Deployment for the city of Stockholm	Open source	ERICSSON
Software	CityPulse PickUp Planner	No	N/A	Planning application	Smart city / Smart transportation	Deployment for the city of Stockholm and Uppsala	Open source	ERICSSON
Software	CityPulse Dynamic Bus Scheduler	No	N/A	Planning application	Smart city / Smart transportation	Deployment for the city of Uppsala	Open source	ERICSSON
Software	Knowledge Acquisition Toolkit	No	N/A	Smart City Data Analytics Algorithms and Software	IoT, smart cities and smart health	Already an open-source tool and used by different research groups;	Open-source	UniS
Software	Ontology Validation Services	No	N/A	Semantic model test and validation	IoT and machine-to-machine	A license is sold to a third-party development for extended development as a commercial tool	Basic version is open-source; extended version is licensed to be commercial	UniS
Software	Social media analyser	No	N/A	Deep learning software for twitter data analytics	IoT and social media analysis	To extend to other cities and integration of cyber-physical and	Open-source	UniS

						social data and create a new toolkit		
Software	Data stream generator	No	N/A	Multivariate data stream generator	IoT and machine learning	A tool for time-series data analysis researchers	Open-source	UniS
General advancement of knowledge	Smart city exhibition CityPulse use cases targeting citizens and SMEs on Trin3.	No	N/A	City Pulse use cases	Smart City	Exhibition/Installation	N/A	AA
General advancement of knowledge /	Knowledge gained about Machine Learning activities	No	N/A	Using the quality monitoring tools developed in CityPulse will result in better planning and utilisation of city services	Smart City	Better public services	N/A	AA
General advancement of knowledge	The relevant smart city use cases from CityPulse will be highlighted within the framework of Open Data DK.	No	N/A	Sharing knowledge through Smart Aarhus, Open Data dk	Smart City	Public communication	N/A	AA
General advancement of knowledge	Engaging with relevant smart city stakeholders on a newly started Smart City communications platform	No	N/A	Smart City Communication Platform	Smart City	Public communication	N/A	AA
Exploitation of results through (social) innovation	The event reporter application will be used to allow citizens to report significant events to be considered by the city via a smart app developed in the project (an already basic version is offered to citizens today: http://www.aarhus.dk/borgertip)	No	N/A	Borgertip APP, Event reporter	Smart City	Better public services	N/A	AA
Software	The Brasov Bus Planner application	No	N/A	Smart planning application	Smart City	Better public services	Open Source	BMA

General advancement of knowledge	To develop the Brasov bus planner into a new app for the cable car transport systems operating the ski slopes of Brasov	No	N/A	Smart planning application	Smart City	New apps / better public services	N/A	BMA
General advancement of knowledge	New sets of relevant urban environment data will be made available on the CKan platform. In order to combine the ideas of the 101 scenarios with the available data, hackathons will be organized starting with the spring of 2017.	No	N/A	CityPulse 101 scenarios / CKan platform	Smart City	Public communication / new apps	N/A	BMA
General advancement of knowledge	One of the exploitation goals is to transfer and adapt the Aarhus demos (travel planner and parking spaces availability) to Brasov, where a couple of projects which will provide the sensors' infrastructure are in preparation	No	N/A	Aarhus City demos	Smart City	New projects / better public services	N/A	BMA
General advancement of knowledge	A selection of the 101 scenarios will be included as project ideas in the local development strategy of Brasov to be implemented during 2016-2023 through the EU structural funds	No	N/A	CityPulse 101 scenarios	Smart City	New projects	N/A	BMA

5 Conclusions

The connected networks of devices and information communication technologies are changing the way that we live and interact with our surrounding world. In the past few years, the amount of collected and published data has been more than ever in human history. The Internet of Things (IoT) and new communication technologies such as 5G are paving the way to collect and process even more data from the physical world. The real world data is becoming more accessible in (near)-real-time and powerful storage and processing hardware and software allow us to handle very large volumes of data. However, we still face the problem of analysing and transforming large volumes of raw data into actionable-information that can be used for decision-making. The Internet of Things and smart cities are full of innovation potential. However, similar to the Web in its early days, they are still in their infancy. The smart city systems are often fragmented and are usually vertical systems aimed for specific environments and applications. Key to overcoming this fragmentation is providing an open ecosystem capable of integrating and processing data from various sources and supporting different applications and services.

CityPulse focused on developing several key enablers for real-world data stream publication and processing, adaptive and scalable data analytics methods, reliability and fault tolerance mechanisms, combined with real world and large-scale demonstrations that are driven by innovative use-case scenarios and new business models. The CityPulse project provides an open and flexible framework for large-scale smart city data analytics and has provided extensive evaluation and analysis of the components and their performance results. This report summarised some of the key outcomes and activities that were conducted during the project. The project website and public report provide more detailed and comprehensive description of the project work and research and development activities and results.

CityPulse is built based on several ideas, models and architectures from previous European Research Framework collaborative projects. This project owes its progress and success to the many individuals and research groups who paved the way for the project team to start with a strong base of collaboration, support and, more importantly, with a clear vision. We would like to thank the European Commission for their support for this project. Our project officer at the European Commission, Mr Eric Gaudillat, was incredibly supportive and encouraged us throughout this project. The reviewers of the project and our extended colleagues in other projects helped us with their feedback, comments and suggestions. Last but not least, our sincere thanks to numerous citizens from different cities across the continent and to our project team and their colleagues who have all contributed to building this project and delivering the results.

References

- [Anantharam et al., 2015] P. Anantharam, P. Barnaghi, K. Thirunarayan, A. P. Sheth, "Extracting City Traffic Events from Social Streams", *ACM Transactions on Intelligent Systems and Technology*, volume 6, issue 4, August 2015.
- [Anantharam et al., 2015-2] P. Anantharam, K. Thirunarayan, S. Marupudi, A. Sheth, T. Banerjee. "Understanding City Traffic Dynamics Utilizing Sensor and Textual Observations ", At The Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16), February 12--17, Phoenix, Arizona, USA. 2015.
- [Barnaghi et al., 2016] P. Barnaghi, A. Sheth, V. Singh, M. Hauswirth, "Physical-Cyber-Social Computing: Looking Back, Looking Forward", *IEEE Internet Computing*, May/June 2015.
- [Barnaghi et al., 2016-2] P. Barnaghi et al., *CityPulse Annual Stakeholder report*, 2016.
- [Farajidavar et al., 2016] N. Farajidavar, S. Kolozali, P. Barnaghi, "Physical-Cyber-Social Similarity Analysis in Smart Cities", (to appear in) the *Proc. of IEEE World Forum on Internet of Things*, December 2016, Reston, VA, USA.
- [Fathy et al., 2016] Y. Fathy, P. Barnaghi, S. Enshaeifar, R. Tafazolli, "A Distributed In-network Indexing Mechanism for the Internet of Things", (to appear in) the *Proc. of IEEE World Forum on Internet of Things*, December 2016, Reston, VA, USA.
- [Flein et al., 2016] *CityPulse Deliverable 6.4, Overall Evaluation Report*, The CityPulse Consortium, 2016.
- [Ganz et al., 2015] F. Ganz, D. Puschmann, P. Barnaghi, F. Carrez, "A Practical Evaluation of Information Processing and Abstraction Techniques for the Internet of Things", *IEEE Internet of Things Journal*, volume: 2, issue: 4, pp. 340 - 354, 2015.
- [Gao et al., 2016] F. Gao, M.I. Ali, E. Curry and A. Mile. "QoS-aware Stream Federation and Optimisation based on Service Composition", *International Journal on Semantic Web and Information Systems (IJSWIS)*, 2016.
- [Gao et al., 2015] F. Gao, M. I. Ali, A. Mileo, "Semantic Discovery and Integration of Urban Data Streams", In the 5th Workshop on Semantics for Smarter Cities, ISWC2015 workshop, 2015.
- [Iggena et al., 2014] T. Iggena, D. Kümper and R. Tönjes, "Kontinuierliche Bewertung von Informationsqualität in Stream-basierten Smart City Architekturen. ITG-Fachbericht-Mobilkommunikation–Technologien und Anwendungen," Osnabrück, May 2014.
- [Iggena et al., 2016] T. Iggena, D. Kümper, M. Fischer, R. Tönjes, "Monitoring and Testing for Reliable Smart City Applications". *ITG-Fachbericht-Mobilkommunikation–Technologien und Anwendungen*", Osnabrück, May 2016.
- [Kolozali et al., 2016] S. Kolozali, D. Puschmann, M. Bermudez-Edo, P. Barnaghi, "On the Effect of Adaptive and Non-Adaptive Analysis of Time-Series Sensory Data", *IEEE Internet of Things Journal*, 2016.
- [Kolozali et al., 2014] S. Kolozali, M. Bermudez-Edo, D. Puschmann, F. Ganz, P. Barnaghi, "A Knowledge-based Approach for Real-Time IoT Data Stream Annotation and Processing", in *Proc. of the 2014 IEEE International Conference on Internet of Things (iThings 2014)*, Taipei, Taiwan, September 2014.

[Kuemper et al., 2016] D. Kuemper, et al., Monitoring Data Stream Reliability in Smart City Environments, IEEE World Forum on Internet of Things (WF-IoT), Reston, USA, 2016

[Mileo et al., 2015] A. Mileo, et al., CityPulse Deliverable D.3.2: "Data Federation and Aggregation in Large-Scale Urban Data Streams", August 2015.

[Nechifor et al., 2014] S. Nechifor, D. Puiu, B. Tarnauca, F. Moldoveanu, "Autonomic aspects of IoT based systems. A logistics domain scheduling example", Workshop on Interoperability and Open-Source Solutions for the Internet of Things, Split, Croatia, 2014.

[Puiu et.al 2016] D. Puiu, P. Barnaghi, R. Toenjes, D. Kuemper, M. I. Ali, A. Mileo, J. X. Parreira, M. Fischer, S. Kolozali, N. Farajidavar, F. Gao, T. Iggena, T. L. Pham, C. S. Nechifor, D. Puschmann, and J. Fernandes, "CityPulse: Large scale data analytics framework for smart cities," IEEE Access, vol. 4, pp. 1086–1108, 2016.

[Puiu et al., 2016-2] D. Puiu et al., CityPulse deliverable D5.3, Smart City Environment User Interfaces, CityPulse consortium, 2016.

[Puiu et al., 2016-3] D. Puiu et al., CityPulse deliverable 7.3, Report on Dissemination Exploitation and Standardisation Activities, CityPulse consortium, 2016.

[Puschmann et al., 2016] D. Puschmann, P. Barnaghi, R., Tafazolli, "Adaptive Clustering for Dynamic IoT Data Streams", (to appear in) IEEE Internet of Things Journal, 2016.

[Puschmann et al., 2016-2] D. Puschmann, P. Barnaghi, R. Tafazolli, "Marginal Distribution Clustering of Multi-variate Streaming IoT Data", (to appear in) the Proc. of IEEE World Forum on Internet of Things, December 2016, Reston, VA, USA.

[Richeard et al., 2000] D. Richeard, P. Hart and D. Stock, Pattern classification, 2nd ed., New York: John Wiley & Sons Inc., 2000.

[Tönjes et al., 2015] R. Tönjes, D. Kümper, and M. Fischer. "Knowledge-based Spatial Reasoning for IoT-Enabled Smart City Applications ", 2015 IEEE International Conference on Data Science and Data Intensive Systems. IEEE, 2015.