Environmental Modelling & Software 102 (2018) 94–114

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

Water utility decision support through the semantic web of things

Shaun Howell*, Yacine Rezgui, Thomas Beach

BRE Institute of Sustainable Engineering, Cardiff University, 52 The Parade, Cardiff, CF24 3AA, UK

ARTICLE INFO

Article history:
Received 11 October 2016
Received in revised form
21 December 2017
Accepted 7 January 2018

Keywords:
Water management
Decision support tool
Interoperability
Big data
Ontology
Semantic web
Internet of things
Smart water networks

ABSTRACT

Urban environments are urgently required to become smarter. However, building advanced applications on the Internet of Things requires seamless interoperability. This paper proposes a water knowledge management platform which extends the Internet of Things towards a Semantic Web of Things, by leveraging the semantic web to address the heterogeneity of web resources. Proof of concept is demonstrated through a decision support tool which leverages both the data-driven and knowledge-based programming interfaces of the platform.

The solution is grounded in a comprehensive ontology and rule base developed with industry experts. This is instantiated from GIS, sensor, and EPANET data for a Welsh pilot. The web service provides discoverability, context, and meaning for the sensor readings stored in a scalable database. An interface displays sensor data and fault inference notifications, leveraging the complementary nature of serving coherent lower and higher-order knowledge.

1. Introduction

In order to tackle sustainability and economic challenges through ICT, urban environments, including the water sector, are undergoing a transformation towards smart systems through the use of web-enabled sensors, analytics software, and decision support tools. Smart water networks have been noted to promote efficacy, efficiency, and resilience of water infrastructure (CTRL+SWAN, n.d.; Mutchek and Williams, 2014; Thompson and Kadiyala, 2014). However, as with fields such as smart grids and smart cities, the application of ICT in the water value chain is restricted due to an inability to share data and knowledge, and hence interoperate, across the people and software components involved (EIP Water, n.d.). This has limited the impact of advanced applications such as optimisation engines, artificial intelligence, and semantic inference. Network operators need modern decision support tools which empower them to make optimal decisions based on extensive data sources and relevant insights, and this interoperability challenge is increasingly pertinent (Curry et al., 2014).

In smart grids, this has been stated by IEEE to occur due to three main issues: lack of machine communication protocols, lack of common data formats and lack of common meaning of exchanged content (IEEE Standards Committee et al., 2011). In the 'smart water' domain, the same core issues have restricted the utility and hence prevalence of ICT penetration. Notably, a recent report from the ICT4Water cluster of EC FP7 projects highlighted the need for standardised models to address the issue of interoperability in the smart water domain (Vamvakeridou-Lyroudia et al., 2015) and specifically indicated the importance of ontologies as a means to maintain semantic clarity and integrate knowledge. This leads to a clear precedent in the smart water domain to develop common communication protocols, data models and semantic vocabularies.

The Internet of Things is addressing the need for interoperable communication protocols, and much progress has been made in the past 5 years, towards enabling device discoverability and message exchange. However, beyond the requirement for applications to receive data, they must be able to consume and utilise it correctly with confidence, which requires a thorough understanding of its context, meaning, and provenance. This requires a robust approach towards semantic interoperability, and achieving this goes beyond the presently observed Internet of Things, through a convergence with the semantic web. This has been termed the Semantic Web of Things (SWoT), with the key difference being a focus on application-layer interoperability, as opposed to protocol-layer interoperability (Calbimonte et al., 2014). SWoT therefore

* Corresponding author.
E-mail addresses: shaunkhowell@gmail.com (S. Howell), RezguiY@cardiff.ac.uk (Y. Rezgui), BeachTH@cardiff.ac.uk (T. Beach).

https://doi.org/10.1016/j.envsoft.2018.01.006
1364-8152/

promises to support advanced applications such as artificial intelligence, where data semantics must be explicitly machine interpretable in order for automated machine-to-machine communication to be sufficient for reliable data utilisation. These explicit semantic statements are typically collated into semantic models, which create a shared understanding of the domain and a shared method of representing data and their meaning. Within this remit, ontologies are the most expressive option, but are also the most complex, in terms of human comprehension and computational complexity. Critically, ontology based models also allow the use of inference to produce new knowledge about a system beyond what has been explicitly stated.

This paper proposes that fostering a Semantic Web of Things will unlock vast potential in the smart water domain, by robustly addressing the interoperability challenge which has been noted as responsible for 40% of the value of IoT systems (Manyika et al., 2015). The paper therefore presents a knowledge-based data and inference platform, and an example decision support tool, which demonstrate the value of the approach. The knowledge management platform couples a scalable time-series database with a semantic knowledge base, empowered by SWRL rules and an accompanying inference engine. This is grounded in a comprehensive domain ontology which describes not only the sensor data and metadata, but also the systemic context of the data, based on a socio-technical system model of the water value chain. This model federates and extends typical GIS schemas, product descriptions, and expert knowledge as an OWL ontology. The ontology is then instantiated through a semi-automatic process, supporting the integration of legacy systems. A key novelty of the approach is the semantic rule base which integrates expert domain knowledge and use case specific heuristic rules with the knowledge-based solution. This is accomplished with real world data from both clean and wastewater networks for a pilot site in Wales, within the context of an EC FP7 research project, entitled ‘Water analytics and Intelligent Sensing for Demand Optimised Management’ (WISDOM).

The value of the hybrid semantic approach is demonstrated through a decision support tool which integrates a semantic rule engine with a comprehensive set of use-case based rules, the platform, and a graphical user interface. The use case of ‘alert impact prediction’ was chosen, whereby a problem such as a blockage occurs in the network, and the decision maker needs to evaluate the impact of the incident on the network and its end-users, within the utility's business context. The application therefore uses semantic inference to predict the affected network components, and allows the user to interrogate the issue by viewing dynamic and static data about the network and problem. This requires an integrated presentation of predicted, static, and dynamic data, in a manner suitable to the expert decision making process. By highlighting the value of leveraging both the knowledge-based, and data-driven programming interfaces of the platform, the decision support tool serves as evidence towards the value of the Semantic Web of Things.

The paper therefore addresses the following primary research question: does a decision support system grounded in IoT and semantic web provide added value over rule based systems currently used in the sector and those observed in literature?

The overarching methodology adopted is that of action research, whereby the researcher works with the stakeholders of a target system to both solve an immediate challenge and generate knowledge from the process and outputs. This involves an iterative learning approach of defining the problem, specifying a proposed solution, building and testing a proposed solution, then reflecting and learning from the iteration. Such an approach allows agility through frequent adaptation, and promotes high-quality outputs through transparency and regular expert review. As such, this paper couples the specification, development, and testing of each of the various components investigated towards answering the stated question.

The rest of the paper proceeds by discussing background concepts and related work in section 2, then section 3 provides an overview of the proposed solution, and section 4 presents the knowledge management platform in detail. The benefit of the approach is demonstrated through an example semantic inference use case built on the platform in section 5, and section 6 presents a decision support tool which leverages the coherent lower and higher-order knowledge served by the platform's programming interface. Section 7 then discusses the work conducted and offers a perspective on the future of knowledge management in smart water.

2. Background

2.1. Smart water and utility decision support tools

The application of ICT and cybernetics principles to the water sector has grown significantly in recent years through the notion of smart water networks (Cahn, 2013; Mutchek and Williams, 2014; Sensus, 2012). These aim to use intelligent sensing (Allen et al., 2013), optimisation (Zhao et al., 2016), and decision support (Schenk, 2010) to operate clean and waste water networks and assets in a more efficient, sustainable and reliable manner. A cluster of European Commission Seventh Framework Programme (EC FP7) research projects, ICT4Water, has been formed to investigate various aspects of this proposition (Vamvakeridou-Lyroudia et al., 2015), and the European Innovation Platform for water (EIP-water) has launched an action group for water monitoring for decision support (EIP Water, n.d.). The smart water networks forum (SWAN) is serving as a nucleus for this trend, and has proposed a framework for smart water networks (Cahn, 2013) consisting of several layers: physical, sensing and control, collection and communication, data management and display, and data fusion and analysis.

One key impact scenario identified by SWAN is intelligent pressure management to reduce leakage and energy consumption whilst improving network resilience. Recent work has demonstrated a 12.5% leakage reduction through intelligent pressure reduction based on an EPANET model (Babel et al., 2009), and this was conferred by another work (Creaco et al., 2016) which minimised energy consumption whilst optimising network pressures. Another work utilised a cloud-based machine learning approach to leakage management (Mounce et al., 2015), and a platform has been developed which aims to integrate ICT with water networks to promote reliable and resilient resource management, whilst reducing energy consumption (Lee et al., 2015), based on a cyber-physical approach. However, implementing these smart water solutions in practice requires pervasive interoperability, as highlighted by the recent SWAN report on communication in smart water (Hauser et al., 2016).

2.2. Supporting smart systems through the internet of things

The role of IoT in smart water has been increasingly noted (Wong and Kerkez, 2016; Robles et al., 2015), resulting in various reference architectures and platforms. It is noteworthy that (Robles et al., 2015) highlights the role of semantics and knowledge-oriented interoperability, although the authors only offer a model of a system's ICT components, as opposed to the underlying socio-technical system, although the paper does stress the value of adding semantics to the work presented.

In the broader smart city field, many examples exist of IoT platforms aiming to coordinate data management (ALMANAC,