Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 94 (2016) 91 – 96

Procedia
Computer Science

The 13th International Conference on Mobile Systems and Pervasive Computing

(MobiSPC 2016)

Towards a Scalable and Optimised Context Broker for Seamless Interoperability within IoT Systems

C B Tokpo Ovengalt[a,b,*], K Djouani[a,b], AM Kurien[a], A Chibani[b]

[a]*Tshwane University of Technology, Private Bag X680 Staatsartillerie Rd, Pretoria 0183, South Africa*
[b]*Université Paris-Est Créteil,61 Avenue du Général Charles de Gaulle, 94010 Créteil Cedex, France.*

Abstract

Billions of devices are now integral part of the day-to-day routine of people but they are sometime incompatible with one another because they employ divergent protocols. The introduction of semantics has been a key enabler towards proper and universal descriptions of connected objects as well as the relationship between them, removing relational ambiguities and improving context-awareness. Virtual objects (VOs) are semantic descriptions of the aggregated native capabilities of physical devices. And as such, they can be readily upgraded over time in order to dynamically suit the processes they are involved with because the context information that is present within them is not necessarily static in time. The original VOs often need to be semantically enriched with more contextual details as they interact with applications and other reasoning engines, where they often evolve into composite virtual objects (cVOs) or rich virtual objects (rVOs) that are more intelligent and interoperable. The complexity of handling large volumes of streaming data increases in the presence of uncertainty caused by incomplete, altered, noisy and/or unstructured information. Context brokering has been a growing trend over the past few years where brokers have been used to discover context and to facilitate processes between service requesting and service providing entities. Several context brokers have been proposed but they do not offer a scalable and optimised brokering algorithm. This necessitates the development of novel ontologies where more functional and non-functional descriptions are used to reduce uncertainties and enrich the content of VOs with the objective of creating a truly automated networking environment.

Peer-review under responsibility of the Conference Program Chairs

Keywords: *Context; Semantics; Broker; Complex Event Processing (CEP); Internet of Things (IoT), Uncertainty.*

---

* Corresponding author. Tel.: +27 764 790 858; fax: +27 12 382 5114.
E-mail address: tokpocb@gmail.com

Peer-review under responsibility of the Conference Program Chairs
doi:10.1016/j.procs.2016.08.016

1. Introduction

Recent technological advances in nanotechnology have led to the embedment of networking capability into sensors and other objects to successfully process and share the data that is produced. The IoT ecosystem is composed of all sorts of physical and virtual things that collaborate to improve several aspects of living with little to no human input. It is of absolute necessity to successfully capture, and process the large streams of data produced by sensors and other things. However, they are issues arising from the lack of interoperability between objects, and uncertainty associated with the reliability and usefulness of the data produced because of the processes involved and the time-varying nature of the information. In this paper, a background of the issues facing the community is presented, the basic composition of VOs is detailed, a review of the state-of-the-art Complex Event Processing (CEP) techniques is presented, and some of the methods used to overcome uncertainty are briefly discussed, a problem facing the research community is formulated, and novel ideas are put forward for further investigations towards a new context broker that offers some degree of scalability, operates optimally and makes decisions in real-time within a fast evolving and heterogeneous IoT environment.

2. Background

The Internet of Things (IoT) is a recently developed concept in which objects that have networking capability interconnect to share the files and data that they produce in order to make decisions in an automated manner. Events are atomic concepts[3] generated by IoT objects such as servers, sensors, smart devices, social media, blogs, on-line services, and so forth[1,7,12]. Events are used to inform a computer program that potentially useful changes in conditions have occurred. Complex events are aggregates or derivations of basic events which are processed to trigger adequate reactions and processes within the systems in which they occur or from peripheral devices. Timely event detection and quality event processing are two critical components to enable seamless networking conditions, especially in autonomous settings. However, the ever expanding device community places a huge demand for scalability on any IoT management system. Billions of triples are generated in short periods of time by the numerous IoT objects, and some of the data produced could be pivotal to the quality of the decisions made by reasoning engines and applications. Additionally, IoT objects often employ different and mismatching protocols, and the vast amount of data that they produce make IoT systems very complex, cumbersome, latent, and occasionally inefficient. This has exacerbated the necessity for quality data processing where semantic clarity, relevancy and context are considered to be the most important ingredients.

For this cause, the World Wide Web Consortium (W3C) developed standards such as RDF, SPARQL and OWL to represent data in unambiguous manner by means of triples that clearly state the relationships between connected nodes, thus reasoning engines can logically process queries at semantic level. In recent years, alliances such as OneM2M and AllSeen were established to enable the exposure of the native capabilities of physical devices and their interfaces through basic VOs that are essentially static in time. However, in real life scenarios, the contents of VOs often need to be enhanced by applications and reasoning engines during processes to be able to comply with some requirements for efficient communication. A VO is a virtual representation of the features, memory, speed, applications, connectivity, etc. of a physical device and its interfaces[8]. As shown in Figure 1, VOs usually have more functionalities than physical objects due to the added capabilities that are derived from surrounding abstract virtual things such as sensors. Whenever necessary, several VOs and abstract things are mashed-up to form new composite virtual objects (or cVOs) with the ability to offer new services that satisfy queries from the end-user.

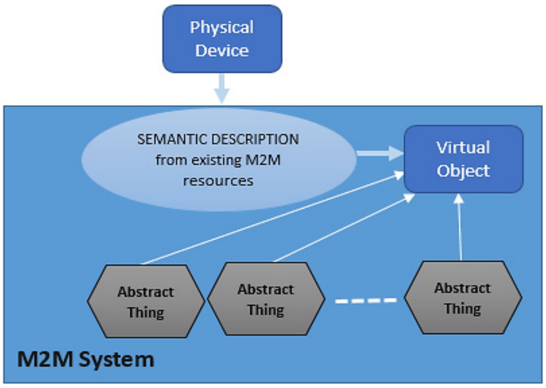

Fig. 1. Virtualization process within the M2M system.

This enrichment concerns mainly context information where functional and non-functional descriptions are added to basic VOs in order to remove ambiguities and promote interoperability. Context is any pertinent information that relates to space, time, ownership, relationship, user-defined preferences, etc. In computer systems, it has become common to use brokers to discover context when direct communication between requester and provider cannot be fully effective due to syntactic and semantic mismatches. A number of context brokers have been developed in recent years but some limitations persist. Only few CEP engines currently address the issue of uncertainty. In the IoT setting, sensors may give faulty reports due to hardware failure or when they encounter a phenomenon they have not been designed for. Handling uncertainty requires the systems to not only identify the sources of imprecisions, but also to either find alternative reliable sources of data, and/or develop an approach that could deal with the missing/incorrect data. A number of approaches have been proposed over the years, including, but not limited to, artificial neural networks, probabilistic logic, fuzzy logic, and Bayesian theory[3].

3. Problem statement

The dynamic nature of IoT ecosystems necessitates the swift detection, capture and real-time processing of the contextual mutations of the content present in basic VOs as they evolved into context rich virtual objects (rVOs) during their interactions with applications and reasoning engines. Furthermore, the uncertainty over the quality of the processed data increases the error rate in the decision making during processes. However, the current context brokers do not have the technical capability scale over huge numbers of triples, nor can they discover and process fast contextual changes in real-time and in the presence of uncertainty. This, in turn, prevents IoT systems from self-adapting in a timely fashion.

4. Preliminary literature

Several complex event processing (CEP) methods were proposed over the years, but most of them do not combine streams with background knowledge nor do they support reasoning tasks over events and that knowledge. However, in IoT settings, background knowledge can be used to enrich the recorded events, detect more events, and make informed and timely recommendations among other things. Semantic Complex Event Processing (SCEP) has been developed to counter some of these shortcomings by adding knowledge management to CEP, but SCEP does not incontrovertibly deal with context-rich information. VIATRA-CEP is a model-based CEP engine, an open-source state-of-the-art processing engine that is developed in the Eclipse platform. It is an integrated modelling environment that employs an explicit runtime model to process events. It integrates the EMF-IncQuery engine to handle the processing of change events in the Eclipse Modelling Framework (EMF) model. One of the main advantages highlighted in this approach is that models are not required to be kept in memory during transformation[11].

VIATRA-CEP uses the VIATRA Event Processing Language (VEPL) for the detection of complex event patterns. VEPL is used to define model transformations and organise them into rules that are set by complex events. However, some of the drawbacks present with VIATRA-CEP include the scalability of a transformation engine, and despite the results obtained in the case study of kinetics, it has not yet been validated in other domains. ETALIS is another open-source framework based on Semantic Web tools that can detect, monitor and prove the relations between events in near real-time, and evaluate domain knowledge on-the-fly[1]. ETALIS is an effective solution for CEP and stream reasoning. ETALIS employs a rule-based language for event called ETALIS Language for Events (ELE). Event Processing (EP) SPARQL[7] is an extension of the SPARQL Protocol and RDF Query Language (SPARQL) on top of which new binary operators such as SEQ, EQUALS, OPTIONALSEQ, and EQUALSOPTIONAL are added to enable event processing and stream reasoning capabilities that are then used to handle real-time Semantic Web applications.

In the ETALIS framework, EP-SPARQL is implemented as an extension to ELE, and Allen's algebra for temporal relationships is used to reason about the intervals at which complex events take place. ETALIS facilitates the management of rules, increases the possibility for sharing among events and intermediate events, and eases the implementation of an operator through "binarization". User-defined pattern rules are compiled into binary rules that are then executed by a standard Prolog system. ETALIS prunes outdated details periodically or according to time windows set by the user. However, like many other reasoning engines, ETALIS uses Semantics that are not clearly and formally characterised in a common language, making it difficult to predict and compare analytically. There are concerns over the ability of ETALIS to scale over large volumes of triples[12]. Furthermore, advanced reasoning features such as non-monotonicity, nondeterminism and model generation are missing.

LARS is a logical framework designed for analysing reasoning over streaming data with different means for time reference and time abstraction[2]. LARS contributes with a rule-based formalism, and different means to refer to or abstract from time, a novel window operator, a model-based semantics, as well as monotonic and non-monotonic semantics. LARS is based on time points but its operators can be used to reason over time intervals as well. LARS aims to be a formal foundation for expressing and analysing different semantic methods for stream processing and reasoning engines[2]. A number of other simulation environments such as Protégé and Virtuoso have been introduced in recent years. However, after using the built-in SPARQL of Protégé 4 and some of the online tutorials, one finds that scalability remains an issue. Virtuoso, on the other hand, allows the processing of millions of triples but there is no support for event polling, a key feature that enables the broker to capture changes in triples in time and in space. The Semantic Sensor Network (SSN) ontology[19] is a key development towards describing the sensing capabilities and properties of sensors by taking into account some environmental factors and deployment conditions. Although this ontology is limited to sensors, it could be extended to include other objects and time properties that will enable the processing of real-time data.

Several techniques have been developed over the years to overcome uncertainty in a number of research areas, including networking, chemistry, robotics, etc. The kernel density estimation (KDE), artificial neural networks (ANN), Bayesian networks (BN), Dempster-Shafer (DS) calculus and fuzzy logic (FL) are among a number of heuristic approaches that have been proposed for reasoning in the presence of uncertainty. Most of these methods are essentially probabilistic, but FL reasons about the interpretation of information. By definition, the KDE[17] is a non-parametric technique in which a smooth approximation of the kernel function is derived from the average of the observed data points. The KDE method handles the variance of the apparent uncertainty by carefully choosing the degree of smoothing which is proportional to the window size. Artificial neural networks (ANN)[16] are algorithms developed to learn, recognise, reason and make decisions based on numerous inputs. Bootstrap and Bayesian statistics are amongst the methods used in ANN algorithms to assess uncertainty.

In Bayesian theory, one seeks to ascertain the correctness of the obtained result via conditional probabilities. The theorem of Bayes states that the posterior probability of the occurrence of event A assuming that event B also occurred, denoted P(A|B), is defined by:

$$P(A|B) = \frac{P(B|A).P(A)}{P(B)} \tag{1}$$

Where P(A) is the *prior* probability that event A occurred, P(B) is the *a priori* non-null probability that event B occurred, and P(B|A) is the likelihood that event B occurs given that event A has occurred. However, the theorem of Bayes cannot be very useful in a hazardous environment because a clear formulation of probabilities is needed, events have to be mutually exclusive and exhaustive, and changes in conditions may be difficult to reflect computationally in real-time. Nevertheless, this theorem forms an important basis to understanding Bayesian Networks (BN)[17] where large numbers of objects and their relationships are graphically represented. BN offer numerous advantages, such as domain knowledge, causal relationships, and, most importantly in this section, handling of incomplete data. The BN approach manages incomplete information by encoding dependencies between sets of input variables. BN uses probabilistic inference for error detection. However, BNs are modelled on the basis of prior knowledge and/or data, while several important probabilities, that are not directly stored in the model, have to be computed.

In Dempster-Shafer (DS) calculus[15], the state of physical systems is formally asserted by means of triples (p,q,r), where p is the probability that the assertion is correct, q is the probability against the assertion, and r is the probability for ‘don’t know’. The representation by means of triples allows for a standard quantification of the degree of uncertainty in terms of probabilities. While in Bayesian theory, $p + q = 1$, the underlying uncertainty triplet r is used in DS analysis because residual ambiguities are implicit in every formal analysis. DS analysis can handle statistical uncertainties by applying DS calculus on a number of DS models of a problem over a defined state space model. DS analysis is valuable to combat uncertainty but a proper formulation of the state space, expertise in DS modelling of problems, and good background knowledge of probabilities are required. Furthermore, real life situations may be very complex and unsuitable for a probabilistic approach because of the amount of computation that would be required and the difficulties in modelling the actual conditions.

Fuzzy logic (FL)[13,14,18] was introduced to account for the imprecise interpretations of crisp values, reasoning about their relationships based on a set of logical rules. Currently, two types of fuzzy sets have attracted the most attention: type-1 and type-2. In type-1 FL, crisp inputs are fuzzified into fuzzy sets that undergo a process of inference based on if-then fuzzy rules. In the fuzzification process, crisp values could have a non-null membership grade $\mu(x) \in [0,1]$ where a degree of ‘1’ indicates full confidence in the exactness of the associated meaning, and likewise, a degree of ‘0’ depicts that the associated meaning is completely inappropriate for the specific crisp value under consideration in one or more MFs. The choice of the number and the shape of the MFs can be helped either by training the system, through stereotyping, by observation, through learning, and so forth. The higher the number of MFs, the more precise and the more complex the system becomes. The shape of MFs may be triangular, trapezoidal, exponential, etc. During the inference process, outputs are mapped based on a set of fuzzy inputs and logical rules.

Type-2 FL is more complex than its type-1 counterpart. Each MF is flanked with an upper and a lower boundary. The area between these two boundaries is known as the footprint of uncertainty (FOU). Some critics of type-1 systems claim that its MFs do not capture the uncertainty well enough. Type-2 on the other hand has MFs that are fuzzy themselves, thereby capturing some amount of imprecision. Another innovation present in type-2 is the type reduction operation where the outcome of the inference process, type-2 fuzzy sets, are reduced into type-1 fuzzy sets before the defuzzification operation. Practical applications have demonstrated that type-2 fuzzy systems handle uncertainty in a better manner than type-1 systems but the former is not always suitable to real-time processing, however efforts are ongoing to make type-2 systems more apt to real-time systems.

5. Future directions of research

The state-of-the-art literature that is presented in the previous section suggests that there is room for contributions towards a scalable and optimised context broker that would be able to adapt to and process contextual changes in real-time despite the imprecise nature of the networking environment. The areas to investigate further include but are not limited to: (1) extending the SPARQL language, (2) elaborating a new reasoning model, (3) developing a new query/answering algorithm, (4) modelling uncertainty, and (5) developing a new simulation environment. Extending the SPARQL language through additional operators has been a key enabler for new technical capabilities that display numerous practical benefits when coupled with a context-based ontology as shown in some of the existing literature[1,4,7]. New functional and non-functional descriptions together with the use of publish/subscribe mechanisms could be key to an efficient monitoring of VOs through semantic annotations enabling the capture of mutations in time, and offer domain-specific support for reasoning over streams. The FIWARE pub/sub generic enablers can be used or adapted for the new context broker as they enable "the publication of context information by context producers, so that published context information become available to context consumers"[10]. The development of a new query/answering algorithm may be necessary to enable the system to respond in timely fashion[12]. For that, some optimisation methods could be explored to ensure that queries are treated in minimum computational time.

A flexible and context-based query model that has an optimum number of steps is desired and/or necessary at the level of the broker. With regards to the use of background knowledge of prior activity and trends within the system, heuristic approaches could also be examined to help reduce processing time through the use of common sense, learning, stereotyping, etc. Heuristics offer the potential to significantly reduce the complexity of the operations of the broker, this is a desirable design property for novel context brokering. Furthermore, heuristics can be coupled with domain-specific background knowledge to deal with uncertainties caused by incomplete, altered, and unstructured information. Critics of heuristics approaches often point to the fact that they do not necessarily yield the most appropriate results, but the trade-off between time saving and finding the best solutions must at least warrant further investigation.

Finally, fuzzy ontologies and reasoning could be used, but again there are concerns over performance[3]. Margara et al.[12] also reckon that "even with efficient incremental algorithm, reasoning may still be too expensive to be performed on-the-fly on streaming data but it may be acceptable in some contexts to trade completeness or precision of reasoning for response time". The use of more general problem-independent approaches such as metaheuristics can also be tested. In truth, some of these methods were already utilised in the literature, but they have not yet been optimally exploited. The development of a new simulation environment cannot be discarded as a mean to achieve scalability. The processing of millions of triples is not a straight forward exercise. Therefore, a combination of existing and/or new techniques need to be considered to improve the state-of-the-art.

6. Conclusion

The basic composition of VOs, the state-of-the-art in CEP techniques, uncertainty handling methods, and some of the problems currently facing the IoT community are presented. Research is ongoing in the field of IoT at local, national, and regional levels around the world. For IoT to become a reality of life, numerous hurdles need to be overcome, and a number of prospective solutions were put forward in this paper for further investigations towards a scalable and optimised context broker that is deployed to achieve seamless interoperability between all the connected devices.

Acknowledgements

This work is partially funded by the French South African Technology Institute (F'SATI) at the Tshwane University of Technology (TUT), in Pretoria, Republic of South Africa.

References

1. Anicic D, Rudolph S, Fodor P, Stojanovic N. Stream Reasoning and Complex Event Processing in ETALIS. *Semantic Web Journal 3(4)*; 2012. p. 397-407.
2. Beck H, Dao-Tran M, Eiter T, Fink M. LARS: A Logic-based Framework for Analysing Reasoning over Streams. *Proceedings of the AAAI conference on Artificial Intelligence*; 2015. p. 1-8.
3. Ye J, Dasiopoulou S, Stevenson G, Meditskos G, Kontopoulos E, Kompatsiaris I, Dobson S. Semantic web technologies in pervasive computing: A survey and research roadmap. *Pervasive and Mobile Computing* (2015), http://dx.doi.org/10.1016/j.pmcj.2014.12.009.
4. Paolucci M, Soudry J, Srinivasan N, Sycara K. A Broker for OWL-S Web Services. *Proceedings of the 1st International Semantic Web Services Symposium*; 2004. p. 1-8.
5. Chen H, Finin T, Joshi A. Semantic Web in the Context Broker Architecture. Proceedings of the 2nd IEEE International Conference on Pervasive Computing and Communication, *IEEE Computer Society*; 2004. p. 1-10.
6. Celik D, Atilla E. A broker-based semantic agent for discovering Semantic Web services through process similarity matching and equivalence considering quality of service. *Science China Information Sciences*. Vol. 56; 2013. p. 1-24.
7. Anicic D, Rudolph S, Fodor P, Stojanovic N. EP-SPARQL: A Unified Language for Event Processing and Stream Reasoning. *Proceedings of WWW*; 2011. p. 635-644.

8. Kim Y, Namgung J, Kim K, Lee GM, Rhee WS. Home Environment Monitoring Service Using Semantic Mash-up. Technical Report - *oneM2M Abstraction & Semantics Capability Enablement.* 2013. p. 1-6.
9. Dey AK, Abowd GD. Towards a Better Understanding of Context and Context-awareness. *CHI 2000 Workshop on the What, Who, When, and How of Context-Awareness*; 2000. p. 304-307.
10. Moltchanov B, Rocha OR. Pub/Sub and Semantic Annotation Enablers for Future Internet in the PPP EU Project Test-Bed. *AAAI 2013 Workshop on Context-Aware System Architectures*; 2013. p. 48-52.
11. David I, Rath I, Varro D. Streaming Model Transformations by Complex Event Processing. *ACM/IEEE 17th International Conference on Model Driven Engineering Languages and Systems;* 2014. p. 1-17.
12. Margara A, Urbani J, van Harmelen F, Bal H. Streaming the Web: Reasoning over Dynamic Data. *Journal of Web Semantics*. Vol. 25; 2014. p. 1-24.
13. Tokpo Ovengalt CB, Djouani K, Kurien A. A Fuzzy Approach to Call Admission Control in LTE Networks. *ANT/SEIT*, Elsevier©; 2014. p. 237-244.
14. Zadeh A. Fuzzy sets. *Information and Control* 8; 1965. p. 338-353.
15. Dempster AP. The Dempster–Shafer calculus for statisticians. *Int. J. Approx. Reason (2007)*. doi:10.1016/j.ijar.2007.03.004.
16. Aires F, Prigent C, Rossow WB. Neural Network Uncertainty Assessment Using Bayesian Statistics with Application to Remote Sensing: 1. Network Weights. *J. Geophys. Res.*; 2004.
17. Taron M, Paragios N, Jolly MP. Modelling shapes with uncertainties: Higher Order Polynomials, Variable Bandwidth Kernels and Non Parametric Density Estimation. *10th IEEE International Conference on Computer Vision*. Vol. 2; 2005. p. 1659-1666.
18. Heckerman D. A Tutorial on Learning with Bayesian Networks. *TechReport MSR-TR-95-06*, Microsoft Research; 1995. p. 1-57.
19. Compton M, Barnaghi PM, Bermudez L, Garcia-Castro R, Corcho O, Cox S, Graybeal J, Hauswirth M, Henson CA, Herzog A, Huang VA, Janowicz K, Kelsey WD, Phuoc DL, Lefort L, Leggieri M, Neuhaus H, Nikolov A, Page KR, Passant A, Sheth AP, Taylor K. The SSN Ontology of the W3C Semantic Sensor Network Incubator Group. *Journal of Web Semantics 17*; 2012. p. 25–32.