Towards Next Generation Building Management Systems

Ovidiu Noran^{1,*}, Ion Sota² and Peter Bernus¹

¹School of Information and Communication Technologies, Griffith University, Queensland, Australia

²Technical University of Constructions Bucharest, Faculty of Installations Engineering, Bucharest, Romania

Abstract. The Internet of Things (IoT) paradigm is gradually finding its way in virtually every industry; however, beyond adding more sensors and measuring and controlling previously inaccessible domains, it is also about transforming ‘legacy’ approaches to control systems, such as those used in Building Management Systems (BMS), by leveraging on the advantages brought by Cyber Physical Systems (CPS). The purpose of this paper is to address several issues gradually emerging in the process of applying the CPS and IoT paradigms to revolutionise BMS. The results of this on-going research aim to help avoid potential pitfalls and provide a sound platform for taking advantage of the benefits brought by this technology in a feasible, effective and controlled manner. More specifically, the paper will address i) the changing meaning of interoperability in the context of the explosion in the number of IoT devices, ii) the need for guidance in adopting sustainable CPS and IoT platforms supporting BMS, based on appropriate non-functional and viable systems principles, iii) emerging issues in the BMS ‘cloudification’ endeavour and iv) the lack of data sources’ correlation resulting in sub-optimal data quality and detail in using Big Data technologies to enable effective analytics for prompt BMS decision-making.

1 Introduction

The Internet of Things (IoT) (“the interconnection via the Internet of computing devices embedded in everyday objects, enabling them to send and receive data” [1]), is finding its way in virtually every industry, including the traditionally local, self-contained Building Management Systems (BMS). Here however, the change brought by the IoT is beyond merely incremental (e.g. more sensors accessing and more actuators controlling new domains). IoT can provide a plethora of opportunities featuring external (e.g. cloud-based) processing and maintenance, augmented intelligence by way of Cyber Physical Systems (CPS) and wider device/service choice avoiding vendor ‘lock-in’ and related high build / maintain costs.

The paper firstly provides a brief review of the benefits brought by IoT in BMS and a description of the most significant problems that unfortunately accompany this development. This is followed by the analysis of each issue found, together with guidance and principles on how to address it and find suitable solutions. This endeavour is performed at a level deemed to be unaffected by the rapid technological changes, so as to achieve stability of the proposed solutions. Finally, the paper summarizes findings and proposes further work.

2 IoT in BMS – Benefits and Issues

According to relevant literature (see e.g. [2, 3]), ‘legacy’, traditionally self-contained and local BMS can be IoT-enhanced by (to list a few):

- Providing access to the wealth of information generated by BMS to external systems which can then in exchange provide efficient analytics [4] for intelligent failure pre-emption and multi-building occupant comfort vs. energy consumption optimisation goals;
- Significantly lowering maintenance and development costs by allowing access to a wider pool of providers [5], changing procurement decision-making and eventually turning product- into service-centricity;
- Using cloud technologies to virtualize the typical server-client BMS legacy architectures [6] and thus help resolve software update, security and vulnerability problems.

Unfortunately, the benefits listed above are typically accompanied by a plethora of issues, which have the potential to prevent achieving the full benefits of the IoT approach. This research has attempted to identify and address the most significant problems, as outlined below:

- The exponential growth in numbers of IoT-enabled devices will determine a high level of heterogeneity of the devices interacting within- and outside a BMS and therefore the commonly understood concept of ‘interoperability’ may need to be revised;
- The adoption of a sustainable and (ideally) *self-evolving* IoT platform for future Smart Buildings is a non-trivial task that requires observing a set of essential principles;
- BMS cloudification is typically not a trivial task; once again, several important architectural principles and requirements need to be observed.

* Corresponding author: O.Noran@griffith.edu.au

- The use of Big Data technologies for effective analytics and use of semi-structured data does not always achieve correlation of data sources or provide information in the quality and detail required for prompt decision-making and thus may lead to false positives and negatives in identifying situations relevant to BMS mission.

These issues must be addressed in order to avoid potential pitfalls and take advantage of the revolutionary benefits brought by the IoT technology in a feasible, effective and controlled manner.

3 Interoperability in IoT-enabled BMS

One of the greatest challenges for the IoT is making the increasingly large number of heterogeneous connected devices exchange the relevant information so they can interoperate. In this context, ‘interoperability’ means that systems must negotiate in real time to achieve meaningful information exchange, even *in the absence of pre-determined assets or agreements for interoperation*. This poses an essential problem with the current definitions of interoperability, based on ‘coexistence awareness’ of the participating systems and agreement of the actors for a given interaction as a result of the mandated interoperation [7].

Unfortunately, the assumption of a pre-existing and agreed-upon interoperability standard is no longer valid in the *ad-hoc* environments created by the large variety of systems involved in ubiquitous computing. The collaboration concept that assumes sharing and a social context may in fact *become a barrier* to interoperability by implying previous agreements between the interoperating systems. Moreover, as the number of connected devices and their technological diversity and complexity increases, it will become more difficult and resource-hungry to reach such pre-agreements.

Therefore, the ‘things’ belonging to the future IoT must be able to receive *ad-hoc* signals and requests from other devices, interpret their meaning and act accordingly. Thus, *each* system must be able to sense, observe, perceive and if necessary, act; thus, interoperability can in fact be considered the *property of a single system*, i.e. an ‘Interoperability as a Property’ (IaaP) paradigm [8].

To exemplify, consider an IoT scenario where a BMS maintenance team with an embedded GPS sensor is deployed within a building complex, moving between BMS service areas (see Fig. 1). Essentially, this sensor (N_1 in the figure) is capable of sensing and perceiving *any* message from its environment (beyond its own).

Furthermore, assume that in the current location of N_1 there are other sensors nodes, observing the local BMS and broadcasting observed data. For example, sensor N_2 is continuously sending message A_{N2} , with CO_2 , air composition or ambient temperature levels in BMS. This message is sensed and observed (O_{N1N2}) by N_1 . In the meantime, the GPS sensor is continuously collecting its own observations (O_{N1N1}). Perception of the service team position, in the context of the air composition of the environment can lead to identifying a

service requirement or even a life-threatening situation for the maintenance / intervention team. In this case N_1 is creating a percept P_1 , based on two observations, namely O_{N1N2} and O_{N1N1} .

Fig. 1. IoT-enabled BMS service / intervention scenario (based on [8])

Based on this perception, N_1 is able to make a decision D_1 , e.g. to send SMS to a command and control centre and/or other teams for assistance or to avoid the danger zone. Hence, N_1 articulates and sends out a message A_{N1} , with request to send SMS with designated content and recipient. Next, there may be a device N_3 (e.g. embedded in the team, ground-based despatch station or other team/s) with SMS sending capability, which observes this message and further acts upon it. This scenario illustrates (in a simplified manner) how a sensor that is able to interoperate with any other sensor/s in the absence of previous protocols.

3.1 IaaP Requirements for IoT-enabled BMS

The scenario shown in Fig. 1 allows to define a set of elementary requirements for the autonomous, intelligent, purposeful and social behaviour of artefacts (or ‘things’ in the IoT meaning) participating in an interoperable environment (e.g. such as a sensor and actuator network engaged in a BMS or belonging to a service team).

Thus, to start with, the artefacts must display *self-awareness* and *environmental awareness*. Self-awareness is related to the capability of the artefact to sense a phenomenon or an event within itself. For example, BMS sensor nodes need to be aware of their available energy- (especially if wireless [9]) and/or functionality level. Environmental awareness is related to the capability of the artefacts to sense a phenomenon or an event and even receive a message from their environment.

Unfortunately, in reality the awareness of nodes is currently only functional in nature and therefore restricted, whereby the sensor is aware only of the environmental features matching its *pre-determined* interest. A similar point can be made related to the capability of the artefact to receive a message of a known format, thus reflecting a limited *functional* (vs. a desired *universal*) environmental awareness.

Perceptivity is another important desired property, whereby the artefacts are able to assign a *meaning* to an

observation. Importantly, observations can occur within the artefacts themselves or in their environment and may also be multi-modal (e.g. temperature, light, pressure, air composition, etc.) and multi-dimensional (e.g. time and location dependent). Perceptivity expresses the achievement of universal awareness by enabling artefacts to observe based on random origins and *interpret* these observations, therefore transforming the physical observations into a meaningful *percept*.

This should further trigger a cognitive process comprising identification, analysis and selection of possible actions. Therefore, artefacts featuring IaaP should possess a third feature, namely *intelligence* - comprising assertion, storing and acquisition of the behaviour patterns based on post-agreements concerning the purposefulness of the performed actions.

Another essential feature of the artefacts featuring IaaP would be *extroversion*, relating to the willingness and capability of the artefact to *articulate* its actions. This would demonstrate the artefacts' interest in the physical and social environment. An associated capability would be 'curiosity', manifested by uttering the request for additional information needed to perform reasoning during the perception and decision processes.

Understandably, the above IaaP requirements imply important associated concerns in regards to ethics, psychology, trust and social aspects to name a few (see e.g. [10]). Treating such aspects in depth is however beyond the scope of this paper.

4 Next Gen BMS: Complexity, Attributes

4.1 The Need to Minimize Complexity

Due to the complex and evolving interrelations among sensor feeds, rising level of intelligence embedded into the measurement and control components [11], the need to often manage clusters of buildings (e.g. within the realm of Smart Cities and Industry 4.0 [12]) and progressively integrating cyber-physical features [13], next generation BMS can be regarded as complex Systems of Systems (SoS). As complex systems cannot be predicted to always satisfy their requirements, BMS SoS complexity should be curbed as much as possible. This could be achieved e.g. by the BMS solution featuring a layered architecture, where the complexity of one layer would not be visible from the layer above, thereby stopping or reducing complexity escalation.

Another effective method of reducing architectural complexity of SoS, using so-called 'Axiomatic Design' [14], has several practical consequences, including the obsolescence of the methodological approach creating a functional specification first, followed by mapping it to a design solution. In this paradigm, even common iterative development- and project knowledge management approaches (such as e.g. agile development) are unsuitable unless they apply the 'zig-zagging' (co-evolution) type of iteration. While complexity reduction techniques would be relatively easy for BMS design teams to acquire, in practice they are often ignored due to historical, rather than technical reasons [15].

4.2 'ilities'

The architectural solution suitable as a foundation for creating cyber-physical systems (CPS) as part of BMS SoS must display a number of systemic properties ('ilities' [16]); moreover, the adopted architecture must ensure that these properties hold true *recursively* for the systems of systems of systems etc.

This requirement originates from the fact that in a SoS the design authority or architect of lower level systems is normally independent of the design authority / architect of the SoS. Thus, the services of the envisaged CPS are to be composed (on a particular SoS level) out of services provided by systems that were independently designed. Therefore, we need extra measures to ensure that at least *service availability, trust, accountability, security, scalability, manageability, longevity, maintainability, reliability, and quality* on each SoS level are achieved and maintained.

A major IoT-based BMS challenge will be innovating using *services upon services* based on core IoT products and services. Thus, a service combination publicly offered may create an *initially* successful and innovative service offering; however, ensuring the above-mentioned systemic properties can be difficult, as business architects must address the design of complex service systems in the context of limited control over underlying 3rd party services (see [17] and Section 5).

The problem is also relevant for providers of the underlying core (e.g., IoT infrastructure) services, as their success depends on the end users' ability to successfully use the infrastructure service over a long period of time. Given that infrastructure providers are few and end users are many, it is in the provider's interest to pro-actively develop *architectural guidelines* to assist successful BMS service composition and thus establish an ecosystem that nurtures innovation [18].

4.3 Viability, Self- and Situation Awareness

A fundamental 'ility' that needs to be singled out in this context is *viability*, which is the property of the system to *self-preserve* and *remain in homeostasis*, but at the same time *co-evolve* with its environment. According to Viable Systems Theory (VSM) [19, 20] a viable system should be composed from equally viable systems; this is essential for the long term survival and success of any large system (such as a complex BMS ecosystem or network) but also for 'virtual enterprises' (VEs) (such as e.g. ephemeral service entities created for specific BMS maintenance projects using a pool of provider and vendor competencies) where the longevity of management and control functions must match the VEs' expected lifetime.

By investigating the management and control functions of viable systems (see Fig.2.) and in light of above considerations, it follows that the SoS in question needs to maintain self-awareness on each level of aggregation, requirement also defined as part of the IaaP paradigm (see Section 3.1). The previously defined self-awareness definition is enhanced here as the ability of the system to

also identify the possibly dynamic relationships between the self and the environment, including other systems. This also includes controlled self-determination and negotiation, i.e., the capability to decide a course of action compatible with internalized principles and understood situation (in line with the percept and IaaP intelligence concepts defined in Section 3.1 and [21]).

Fig.2. Recursive Management and Control functions of a viable system (based on [22], combining Beer's VSM [19] and GRAI Grid [23] views)

The system of interest (SoS) should have functions to provide services (control, maintenance, upgrade etc.) but also must have functions to monitor the ability of the system to perform these functions presently *and in the future* (such as through monitoring the performance of the self and of the environment). This is not usually the case with lower level granularity systems, and can cause unpredictability and brittleness on the SoS level.

Clearly, self-awareness and viability are necessary to ensure system homeostasis (maintaining all necessary 'ilities' discussed above), but when deemed necessary also to self-evolve. For this reason, the SoS in question (e.g. the CPS-BMS SoS) can be regarded as a hybrid (human-machine) system, as opposed to the traditional systems engineering view dividing the human and the machine early in the design, thus considering separately the organization of the system (which is automated) and that of the humans (who are the 'users'). The advantage of this hybrid approach is twofold: a) lack of constrain by design to only implement management functions that can be automated at any one time, and b) no separation of the system along a boundary between two parts with substantial coupling (human-to-machine).

This approach is equivalent to architecting the CPS-BMS SoS as a *multi-agent* system; however, while the multi-agent systems analysis is focused on fully automated individual- and cooperative agents (e.g.. robot swarms), this approach chooses to not constrain these agents *a priori* to full automation, in order to avoid a pre-conceived implementation decision. The removal of this constraint allows for an independent evolution of agents that changes the level of automation in time but still preserves an architectural identity and with that, longevity. Thus, by default the system always has a

complete scope and *all necessary levels of management and control*, but as the system evolves, the level of automation changes (while preserving system identity).

4.3.1 Self-awareness example: real time and operational level

On the real time level, it is necessary for the CPS-BMS SoS to e.g. identify faults (of the self, *or of external services*, e.g. for maintaining security [24, 25]), identify cyber-attacks, or any other situation that demands action that flows *in sync* with the events of the process. This needs the constant *timely* evaluation of data streams (see Section 6) and their interpretation (into *percepts* as shown in Section 3.1) to create situation awareness and enable intelligent selection of the appropriate response. Due to human limits to promptly process and act in the context of exponential data and complexity, next generation BMS self-aware behaviour on the real time- and operational levels should be *highly automated* (a specific feature of cyber-physical systems, including those BMS-bound).

For example, it is essential that even at the operational level (e.g. executable code) some level of situation awareness (as an internal control) is present. If this is not true, an 'atomic service' is open to compromise and may be executed in an improper fashion by an illegitimate 3rd party. Emerging CPS-enabled BMS do not appear to possess this ability [26]; therefore, cyber-attacks that enter on a very low level may remain unnoticed. In the emerging complex BMS environment this is a significant threat that could result in nuisance, disruption, life safety / panic, ransom requests, privacy violations, override of security and other vital building management parameters [27]. Building situational awareness (thus the ability to respond to known and unknown situations based on available and newly arriving data in real time) is a potential enabler to fend off emerging cyber-threats.

4.3.2 Situation Awareness

A recurring theme in the discussion above is situation awareness (coupled with fast decision making ability) for effective and efficient action. The conditions of this to materialise are facilitated by the technological affordances of IoT (large number of intelligent sensors), machine learning / pattern recognition algorithms, and various data analytics techniques. However, highly automated situation awareness and supporting decision making requires a new form of intelligence, i.e. situated reasoning [28, 29], in order to create a *fast and continuous narrative* of the facts uncovered by the variety of data sources. Situated reasoning and context level data analysis should go together: the analytics of large amounts of data can facilitate correct situation identification, while situated reasoning can identify the *need* for data that are not available at present, but *could be used* for situation disambiguation before a correct decision can be taken [30]. This aspect is described in more detail in Section 5.

5 CPS-BMS Cloudification

Cloud computing holds the promise to allow enterprises deploy and operate applications faster, reduce maintenance and improve manageability [31, 32]. Figure 3 presents a typical cloud architecture, which, while following the accepted National Institute of Science And Technology (NIST) terminology (Liu et al., 2011), considers the architecture from a *functional* point of view only (i.e., leaving it open on how service, service management and other related functions are distributed among participants). As can be seen, in Fig. 3 Information Services are also treated separately, given new forms of services that are based on various other information sources other than databases (such as e.g. streaming data from sensors within a BMS).

Fig.3. Possible BMS Cloud Architecture based on NIST Cloud Computing Reference Architecture [33].

An increasing number of problems are surfacing in the ‘cloudification’ of next generation BMS, partly due to the expected inherent complexity of- and turbulence created by any major change; however, there are also specific problems originating in the very nature of cloud computing and importantly, in the fact that both the emerging model of CPS-BMS, as well as the cloud computing business model and technology themselves evolve during and after the ‘cloudification’ project(s).

A first major hurdle is the extent of cloudification; thus, organisations must apply due diligence when selecting and moving functionality to the cloud, as cost and productivity advantages also bring potential drawbacks in risk and liability [34]. For example, although government organizations (e.g. Defence-related) may be discouraged to outsource BMS services to a public cloud, appropriate risk mitigation could allow a partial deployment in the cloud [35].

A second important challenge increasing the complexity of the cloudification endeavour is service recursion, where services may call upon other services, such as e.g. in an ‘intercloud’ architecture [36]. For example, the next generation BMS may need services such as security, maintenance, upgrade etc. which may be best (e.g. cost-effective) offered as external services [37]. This situation raises several important questions such as: Who is responsible for reliability if a service

fails due to other services it depends on? Can a service be guaranteed if it is integrated with / depends on others? If so, who is the responsible / guarantor entity? Therefore, it is important that the degree of recursion is properly understood and therefore, adequate queries are raised by the acquisition panel of the to-be-cloudified business to the cloud computing solution provider.

A main driver for cloudification is the promise of lowered costs; therefore, realising if appropriate savings are indeed achieved in the long run is an important issue (considering also that business needs will change). Thus, the total cost of ownership (including initial migration and deployment to the cloud, operation, continual development and decommissioning / migration) can escalate if the cloud pricing model is not understood and strategically assured by the user. The use of cloud pricing frameworks (see e.g. [38]) is useful in this regard so as to understand the options; however, this should not be the only factor and as previously stated, first of all it should be decided why cloudification is necessary or desirable, what is to be cloudified, and to what extent.

A heterogeneous cloudification solution (using several providers) could be cheaper and a better fit for purpose as various providers offer different coverage of specific services and thus best prices may be negotiated for each application type. This option would require more varied in-house competencies compared to relying on a single cloud service provider; however, the latter option has drawbacks such as potential lock-in, or high exit cost should a migration be necessary.

Zardari et al. [39] argue that analysing Service Level Agreements (SLA) of cloud providers and matching them against the user requirements can reveal potential violations of important principles, or conflicts and risks; the above discussion has revealed however, a much broader spectrum of cloudification challenges. Due to their intertwined character, these challenges have to be addressed in a more holistic manner, based on the entire set of applicable quality of service (or ‘architecturally significant’ [40]) requirements. Thus, the question is which non-functional systemic requirements, or ‘ilities’ (see [16] and Section 4.2) are affected by the various cloudification solutions, how and to what extent.

Thus, it appears that two main current problems in achieving a successful cloudification are:

- 1) Cloudification cannot be just done off the shelf - the business needs to transform to some extent (while still operating as it cannot afford to stop) so as to minimize turbulence and best take advantage of the service structure offered by the cloud, and

- 2) Even if 1) is accomplished, how can the end user avoid the undesirable side effects of moving to the cloud? The current lack of emphasis on the interactions between the various entities inhabiting the layers created by 1) and the interaction among the life cycle phases of these entities carries the risk of creating sub-standard solutions that suffer from multiple systemic aspects in unanticipated ways. Therefore, cloudification requires a *holistic* approach that considers BMS mission fulfilment as well as management and control, to only name a few

viewpoints. See [22] for more details on suitable architectural frameworks for business cloudification.

6 Efficient Data Analytics for CPS-BMS

An important component of BMS is about reasoning and making decisions (e.g. transmitted to actuators) based on interpreted and hopefully *understood* (abstracted at a ‘meta’ level) data collected from sensors (see also Section 3). As many legacy BMS are rarely updated and maintained, they are a source of rather anecdotal and low quality metadata [13]. Various solutions have been proposed to improve this situation (see e.g. [13, 41]); however, to the best knowledge of the authors, none of these tackles the problems faced by current technologies like ‘big data’ in reasoning under uncertainty.

6.1 An Effective Decisions Model for CPS-BMS

In a complex SoS such as a CPS-BMS, the tasks that appear in each type and level of decision-making and the

feedback that can be used to inform the filters used to selectively observe reality, may be studied using models that explain how successful decisions are made. One such model is the so-called Observe, Orient, Decide and Act (OODA) Loop devised by John Boyd [42].

Despite some opinions of the contrary [43], OODA is not a strict loop, due to the feedback links inside the high level ‘loop-like’ structure that are responsible for learning and for decisions about the kind of filters necessary. Thus, in fact it is actually an activity network featuring rich information flows among the OODA activities and the environment, a very important aspect in view of the chosen BMS decisional scope and the current Big Data approaches tendency to disregard the context of collected data [44, 45]. Thus, OODA can highlight potential caveats and development directions for big data methodology applied to decision support. The OODA tasks are as follows (see Fig. 4):

- Observe (*selectively perceive*) data – measurement, sensors, real-time data streams using existing sensors);

Fig.4. Extended OODA Loop as an activity network (based on [46] and [47]) featuring added Learning and Narratives Loops

- Orient (*recognise and become aware of the situation* based on patterns in data using analytics and producing a narrative to what is actually happening);
- Decide (*retrieve existing-, or design / plan new patterns* of behaviour);
- Act (execute behaviour, then observe outcome, etc.).

Note that although progress in sensor networks and the IoT can provide data on a massive scale, it is impossible to observe *everything*; therefore, there is no certainty that what is observed is relevant and can be analysed to obtain all information necessary or useful *situational awareness* [48] and effective decision and action. Only through ‘post-mortem’ *learning* a BMS can acquire timely and effective analytics supporting the above mentioned capability. Importantly, this learning *is in itself another OODA loop* (see Learning Loop in Fig. 4) featuring various questions: i) what to observe, ii) how to orient to become situation-aware and iii) what is guiding the decision about ‘what to do’ (within constraints, decision variables and possible actions).

Essentially, such strategic self-reflection compares the current capabilities of the BMS to desired future capabilities, enabling management / control to decide whether the change affects system capabilities (including decision making), identity (re-missioning), or both.

6.2 Big Data Decision Support Consequences

The above analysis implies that ‘big data’ (i.e. the collective technologies and methods of data analysis and predictive analytics) does have the *potential* to enable situational awareness (a condition of successful action) by delivering previously unavailable domain-level facts and patterns relevant for decision-making. However, this data needs to be *interpreted*, which calls for a *theory of situations* resulting in a *narrative* of what is being identified or predicted. Without such a narrative, there is no true situational awareness or trust in the system, which can seriously limit the chances of effective action.

Therefore, having the ability to gather, store and analyse large amounts of data using only algorithms is not a guarantee that the patterns thus found in data can be turned into useful and trustworthy information that forms the basis of effective decision-making, followed by appropriate action/s leading to measurable success.

Importantly, this is also true the other way around: when interpreting available data, there may be multiple fitting narratives and it is difficult to choose the 'correct' one. Here, reasoning with incomplete information could help articulate a need for new data (or new types thereof) that could resolve the ambiguity.

Based on the above arguments, the authors argue that decision-making based on data warehousing and/or using 'big data' requires the collection of a *second level* of data. This 'second level' does not refer to particular facts, but rather underpins the creation of an *inventory of situation types*, with facts that *must* be true, facts that *must be not* true, as well as constraints and rules of corresponding *causes* and *effects*. These situation types can be considered models of the domain that can be matched against findings on the observed data level.

Note that due to the ever-changing nature of the Universe of Discourse, one should not aim to design and construct a facility that relies on a completely predefined ontology of situation types. Rather, there is a need for a capability to continuously improve and extend this type of knowledge, including learning of new types (i.e. not a specialisation of some known type). This is required in order to ensure that the 'world of situations' remains open, as described by Goranson and Cardier [29].

In order to achieve adequate situation awareness for effective decision making, collected data needs to be filtered based on relevance [49], dictated by the possible situations of interest. However, as the current situation is typically ambiguous and changes, one will have to maintain a set of dynamic narratives that will continually adjust data needs as well as what needs to be filtered out or kept. This constitutes yet another OODA loop, applied to the set of narratives assisting in the interpretation of data for decision making (see Narratives Loop in 4).

7 Conclusions

The advent of IoT technologies holds the potential to revolutionise the legacy methods used in BMS and support the application of CPS paradigms to BMS. However, the promised advantages are accompanied by some essential caveats. This paper has attempted to describe and propose ways of dealing with these issues, starting with the predicted need for a new type of interoperability (as a property, IaAP) brought by IoT and associated requirements for achieving it.

Further on, the paper has advocated the necessity for the IoT platform adoption strategy to consider essential enablers such as self-awareness, 'ilities' such as viability, availability, reliability while also providing design guidelines for integrating CPS into the next generation BMS. This was followed by a description of problems encountered and strategies for possible solutions in the emerging trend of moving the BMS in

the 'cloud'. The authors have also proposed an enhanced OODA loop-based approach featuring a second level of situated logic for the efficient use of data analytics, so as to effectively support the dynamic configuration and reconfiguration of the CPS-BMS SoS for resilience, efficiency and other desired systemic properties.

Future work will focus on developing 'proofs of concept' for the strategies and high-level solutions proposed, augmented by several real-world case studies.

References

1. Oxford Dictionaries. *Internet of Things Definition*. https://en.oxforddictionaries.com/definition/internet_of_things (2018) [cited 2019 Jan.]
2. M. Ernst. *IoT For Smart Buildings Isn't What You Think It Is*. <https://medium.com/iotforall/iot-for-smart-buildings-isnt-what-you-think-it-is-bc4019270a47>, (2018) [cited 2018 Dec.]
3. B.S. Brad and M.M. Murar, Smart Buildings using IoT Technologies. Construction of Unique Buildings and Structures. **5**(20): p. 15-27, (2014)
4. J. Petze. *Using Data to Improve Facility Operations*. <https://newdeal.blog/using-data-to-improve-facility-operations-78bb1d1b0580>, (2017) [cited 2018 Dec.]
5. T. Kannegieter. *The IoT and Building Management* <https://www.ecdonline.com.au/content/article/the-iot-and-building-management>, (2018)
6. S. Caluianu and F.A. Hebean, Cloud Computing and Internet of Things Concepts applied on Buildings Data Analysis. Mathematical Modelling in Civil Engineering. **13**(4): p. 39-49, (2017)
7. S. Soursos, et al., *Towards the cross-domain interoperability of IoT platforms*, in *EU Conf. on Networks and Communications (EuCNC)*: Athens, Greece. p. 398-402, (2016).
8. O. Noran and M. Zdravković, *Interoperability as a Property: Enabling an Agile Disaster Management Approach*, in *ICIST 2014*: Kopaonik, Serbia, (2014)
9. Q. Chi, et al., A Reconfigurable Smart Sensor Interface for Industrial WSN in IoT Environment. IEEE Trans Ind. Inf., **10**(2): p. 1417-1425, (2014)
10. D. Sikeridis, et al., *Socio-Physical Energy-Efficient Operation in the Internet of Multipurpose Things*, in *2018 IEEE Int'l Conference on Communications (ICC)*: Kansas City, MO. p. 1-7, (2018).
11. M. Manic, et al., Building Energy Management Systems: The Age of Intelligent and Adaptive Building. IEEE Industrial Electronics Magazine. **10**(1): p. 25-39, (2016)
12. J. Ploennigs, A. Ba, and M. Barry, Materializing the Promises of Cognitive IoT: How Cognitive Buildings are Shaping the Way. IEEE Internet of Things J. **5**(4): p. 2367-2374, (2018)
13. J. Fürst, et al., *Crowd-sourced BMS Point Matching and Metadata Maintenance with Babel*, in *PerCom Workshops*. p. 1-6, (2016).

14. N. Suh, *Complexity: Theory and Applications*. Oxford University Press, (2005).
15. I. Gorzeń-Mitka and M. Okręglińska, Managing Complexity: Current Strategies and Approaches. *Procedia Econ & Fin.* **27**: p. 438-444, (2015)
16. O.L. de Weck, *Life-Cycle Properties of Engineering Systems: The Ilities*, in *En. Sys.*, O. de Weck, D. Roos, and C. Magee, (Eds.). p. 65-96, (2011).
17. J.R. Rabelo, O. Noran, and P. Bernus, *Towards the Next Gen Service Oriented Enterprise Architecture*, in *Proc. IEEE 19th Int. Enterprise Distributed Object Computing Workshop*. p. 91-100, (2015).
18. R.J. Rabelo, P. Bernus, and D. Romero, Innovation Ecosystems: A Collaborative Networks Perspective. Risks and Resilience of Collaborative Networks. *IFIP AICT* **463**: p. 323-336, (2015)
19. S. Beer, *Brain of the firm*. London: Allan Lane Penguin Press, (1972).
20. P. Hoverstadt, *The Fractal Organization: Creating sustainable organizations with the Viable System Model*. Hoboken, N.J.: Wiley, (2008).
21. P. Turner, P. Bernus, and O. Noran, Enterprise Thinking for Self-aware Systems. *IFAC Papers OnLine*. **51**(11): p. 782-289, (2018)
22. O. Noran and P. Bernus, *Business Cloudification: An Enterprise Architecture Perspective.*, in *Procs of ICEIS2017* J. Filipe, et al., (Eds.). ScitePress: Porto, Portugal. p. 353-360, (2017).
23. G. Doumeingts, B. Vallespir, and D. Chen, *GRAI Grid Decisional Modelling*, in *Handbook on Architectures of IS* P. Bernus, K. Mertins, and G. Schmidt, (Eds.). Springer Verl.: p. 313-339, (1998).
24. T. Olavsrud. *11 Steps Attackers Took to Crack Target*. <http://www.cio.com/article/2600345/security0/11-steps-attackers-took-to-crack-target.html>, (2014) [cited 2019 Jan]
25. K. Zetter. *Researchers Hack Building Control System at Google Australia Office*. *Wired.com* <https://www.wired.com/2013/05/googles-control-system-hacked>, (2013) [cited 2018 Dec]
26. K. Paridari, et al. *Cyber-Physical-Security Framework for Building Energy Management System*. in *2016 ACM/IEEE 7th (ICCPs)*. (2016).
27. C. Grundy, Cybersecurity in the built environment: Can your building be hacked? *Corporate Real Estate Journal*. **7**(1): p. 39-50, (2017)
28. K.J. Devlin, *Logic and Information*. Cambridge University Press, (1995).
29. T. Goranson and B. Cardier, A two-sorted logic for structurally modelling systems. *Progress in Biophysics / Molecular Bio*. **113**: p. 141-178, (2013)
30. P. Bernus and O. Noran, Data Rich – But Information Poor. *IFIP Advances in ICT*. **506**: p. 206-214, (2017)
31. M. Hirzalla, *Realizing Business Agility Requirements through SOA and Cloud Computing*, in *18th IEEE Int'l Req Eng Conf*. p. 379-380, (2010).
32. M. Sawas and M. Watfa, The impact of cloud computing on information systems agility. *Australasian J of Info. Sys.* **19**: p. 97-112, (2015)
33. F. Liu, et al., *NIST Cloud Computing Reference Architecture*. *NIST SP 500-292*, Gaithersburg, MD: NIST IT Laboratory, (2011).
34. E. Cayirci, et al., A risk assessment model for selecting Cloud Service Providers. *J. Cloud Computing*. **5**(14), (2016)
35. W. Jansen and T. Grance, Guidelines on security and privacy in public cloud computing. NIST special publication 800-144. (2011)
36. M. Morrow, et al., *Blueprint for the Intercloud - Protocols and Formats for Cloud Computing Interoperability*, in *Int'l Conference on Internet and Web Applications and Services*. p. 328-336, (2009).
37. F. Aulkemeier, et al., A pluggable service platform architecture for e-commerce. *J. Inf Syst E-Bus Management*. **14**: p. 469-489, (2016)
38. G. Laatikainen, A. Ojala, and O. Mazhelis, *Cloud Services Pricing Models*, in *Software Business*. G. Herzwurm and T. Margaria, (Eds.). Springer: Berlin Heidelberg, (2013).
39. S. Zardari, F. Faniyi, and R. Bahsoon, *Using obstacles for systematically modeling, analysing, and mitigating risks in cloud adoption*, in *Aligning Ent. Sys. & Sw Arch*. IGI Global. p. 275-296, (2012).
40. L. Chen, M. Ali Babar, and B. Nuseibeh, Characterizing Architecturally Significant Requirements. *IEEE S'ware*. **30**(2): p. 38-45, (2013)
41. J. Gao, J. Ploennigs, and M. Berges, *A data-driven meta-data inference framework for building automation systems*", in *BuildSys - 2nd ACM Conf. on Embedded Sensing Systems...* p. 23-32, (2015).
42. F. Osinga, *Science, strategy and war: The strategic theory John Boyd.*, London, UK: Routledge, (2006).
43. K. Benson and S. Rotkoff, Goodbye, OODA loop: A complex world demands a different kind of decision-making. *Armed Forces J.* **149**(3): p. 26-28, (2011)
44. B. Marr, *Big Data: Using Smart Big Data, Analytics and Metrics to Make Better Decisions and Improve Performance* Wiley & Sons, (2015).
45. R. Corrigan, *Digital Decision Making: Back to the Future*. Springer Verlag, (2007).
46. O. Noran and P. Bernus, *Improving Digital Decision Making Through Situational Awareness.*, in *Info. Systems Development: Designing Digitalization* B. Andersson, et al., (Eds.): Lund. Sweden, (2018).
47. D.S. Fadok, J. Boyd, and J. Warden, *Air Power's Quest for Strategic Paralysis*, ed. Maxwell Air Force Base AL. Air University Press, (1995).
48. V. Lenders, A. Tanner, and A. Blarer, Gaining an Edge in Cyberspace with Advanced Situational Awareness. *IEEE Sec & Priv* **13**(2): p. 65-74, (2015)
49. M. Li, et al., *Big Data-driven Technology Innovation: Concept and Key Problems*, in *Procs of WHICEB 2017*. AIS Electronic Library, (2017).