Olympus: The Cloud of Sensors

Igor L. Santos, Luci Pirmez, and Flavia C. Delicato, Universidade Federal do Rio de Janeiro
Samee U. Khan, North Dakota State University
Albert Y. Zomaya, University of Sydney

A decentralized wireless sensor and actuator network (WSAN) virtualization model leverages the cloud of sensors paradigm to make the best use of the cloud and physical WSAN environments.

In the last few years, we've witnessed the emergence of a new paradigm, the cloud of things (CoT),[1] a combination of cloud computing[2] and the Internet of Things (IoT).[3] A cloud is a large-scale (ideally unlimited) set of user-friendly virtualized computing resources that can be dynamically reconfigured to serve a variable load, seeking optimum resource utilization.[2] The IoT paradigm envisions a global network infrastructure linking a wide variety of physical and virtual devices—the "smart things" that provide identification, data processing, sensing, and connection capabilities to support the development of cooperative services and applications.[3] Essentially, in the CoT paradigm, the cloud acts as an intermediate layer between smart things and applications. Such an intermediate layer hides the complexity of smart things necessary to implement applications.[1] The IoT can benefit from the cloud's virtually unlimited resources to implement service management and composition for utilizing smart things and the data they produce,[1] whereas the cloud can benefit from the IoT by extending its scope to deal with real-world objects (smart things) in a distributed and dynamic way.

Smart sensors play an important role in the CoT paradigm.[1] These sensors are tiny battery-powered devices endowed with processing, storage, sensing, actuation, and wireless communication capabilities. Such capabilities enable smart sensors to be grouped together to monitor variables, forming a wireless sensor and actuator network (WSAN).[3] WSANs include sink nodes, which are nodes at the edge of the WSAN that don't have the computational, energy, or communication resource constraints of smart sensors. Sink nodes serve as entry points for application requests, as well as points for collecting data from the smart sensors. In WSANs, the data acquired by the smart sensors can be processed and interpreted locally and/or sent to one or more of the sink nodes. Actuators can perform actions on the physical environments in response to the smart sensors' decisions.

Given these smart sensor capabilities, WSANs show an advantage within the CoT paradigm in their support of the development of myriad novel coop-

IEEE CLOUD COMPUTING | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | Qmags THE WORLD'S NEWSSTAND®

erative services and applications. However, WSAN devices are also less heterogeneous and mobile than some devices typical of IoT, such as wearable sensors or field operation devices. They're also more resource constrained, specifically in terms of the available energy for operation.[4] Therefore, the potential advantages of WSANs along with smart sensors' specific features have motivated the emergence of the *cloud of sensors* (CoS) paradigm as a type of ecosystem within the broader domain of CoT.[4] A CoS is composed of virtual nodes built on top of physical WSAN nodes and provides support to several applications that might in turn require access to functionalities at the infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS) levels. Application owners can automatically and dynamically provision such virtual nodes on the basis of application requirements. In this sense, CoS infrastructures are built on the concept of WSAN virtualization,[5] which is expected to provide a clean decoupling between services and infrastructure.

According to Sanjay Madria and his colleagues, such CoS infrastructures, built on the concept of WSAN virtualization, have several advantages in terms of WSAN node management, the possibility of sharing data captured by physical sensors among multiple users, and transparency, from the users' viewpoint, regarding the type, distribution, and location of physical sensors in use by an application.[4] Moreover, because the CoS infrastructure follows an open, flexible, and reconfigurable design concept,[4] it can support applications requiring large-scale WSAN deployments, such as precision agriculture and structural health monitoring (see the related work sidebar).[6] Despite these advantages, however, there's still at least one important challenge to deal with in their design. This challenge pertains to the development of a model for WSAN virtualization that simultaneously meets the requirements of several applications, while dealing with the resource-constrained nature of WSANs and prolonging their lifespan.

Here, we discuss this challenge of WSAN virtualization and the drawbacks of existing WSAN virtualization approaches. In addition we introduce Olympus, our WSAN virtualization model.

The Challenge of WSAN Virtualization

Several proposed CoS infrastructures consider physical sensors as passive devices able to provide data to the closest sink node, which forwards such data to a (often) single database stored in the cloud.[2,4,7] Being fully inside the cloud, WSAN virtualization takes place based simply on processing/correlating data stored in this database, and thus in a centralized manner. Such a model is traditionally known as *sensing as a service*.[4] Moreover, CoS infrastructures are usually based on publish/subscribe mechanisms, where each physical sensor publishes the sensed data and metadata (comprising sensor types, locations, and other useful descriptive information), and applications subscribe to the published sensor data.[2] Each application subscription to a published set of sensor data results in the instantiation of new virtual sensors or the reuse of existing ones. Consequently, the instances of virtual sensors are created and exist only inside the cloud, based on the (possibly correlated) data provided by the existing physical sensors connected to the CoS. These centralized CoS infrastructures demand transmission of data to a sink node connected to the WSAN. The communication overhead caused by such an approach is aggravated when large-scale physical deployments are used to increase the frequency of simultaneous transmissions. This communication overhead results in drawbacks that hinder centralized CoS infrastructures from achieving better results when used as solutions to the challenge of the WSAN virtualization.

> In Olympus, application decision processes are performed partly within physical sensors and partly within the cloud.

There are several drawbacks to the centralized approach. The first, communication overhead, compromises the WSAN lifespan, because the nodes of the WSAN have limited energy resources.[4,6,8] An effective solution for maximizing the system lifespan is to process and reduce the sensed data locally, within the physical WSAN. Such data reduction consists of decreasing the amount of data (which reduces the corresponding transmissions) used by applications to make decisions. A possible approach for data reduction is to use *information fusion*,[8] which consists of transforming/joining (fusing) two or more pieces of information (data) from different sources, resulting in other information. In this approach, it's possible to consider virtual sensors

RELATED WORK ON WSANS IN CLOUD OF SENSOR ARCHITECTURES

Dung Phan and his colleagues proposed the sensor-cloud integration platform as a service, a cloud-integrated wireless sensor and actuator network (WSAN) architecture.[1] SC-iPaaS virtually represents physical system components within clouds by accessing the components through virtual representations and processing/managing the data collected within the clouds. A push-pull communication scheme is adopted among the three layers of the SC-IPaaS, in which the lower layer comprises the physical sensors, the middle layer comprises the sink nodes, and the higher layer comprises the cloud applications. In push communications, the physical nodes periodically transmit data to the sinks that forward such data to the virtual sensors. In pull communications, the applications request data that is unavailable to a virtual sensor. Phan and his colleagues also formulate a problem seeking optimal data transmission rates for the individual nodes within the three-tier push-pull communication scheme.

Sanjay Madria and his colleagues proposed a centralized cloud of sensors (CoS) architecture in which a virtual sensor is defined as an emulation of a physical sensor and obtains data from the underlying physical sensors.[2] In such an architecture, virtual sensors are implemented as an image within the software of the corresponding physical sensors. That is, the virtualization software is partly within the cloud and partly within the physical sensors. However, the part within the sensors is used only for communicating sensory data and the metadata to the cloud, resulting in a centralized WSAN virtualization model.

Kyu Hyung Kim and his colleagues proposed a centralized CoS infrastructure for supporting agriculture applications including air/soil monitoring, crop control, growth status monitoring, and surveillance, all of which require managing large-scale WSANs.[3] The proposed CoS infrastructure includes a service layer, in which several applications (built from templates) run. It also includes a virtual layer, in which virtual sensors, actuators, and gateways are provisioned to applications, and a physical layer comprising physical WSANs. The authors proposed solutions for the efficient management of sensors, real-time processing and storage of WSAN data, and the provision of various services such as efficient data communication over physical WSANs, multipath source routing for guaranteeing reliability and fault tolerance, and hierarchical routing to solve scalability issues.

Imran Khan and his colleagues identified two approaches to allow multiple applications to access WSAN resources.[4] Centralized WSAN virtualization models use the first approach, which is to allow multiple applications to share the data gathered from a WSAN. In this approach, a sink/gateway node collects all the data from the WSAN and shares it among multiple users. Decentralized WSAN virtualization models use the second approach, which is to use the capabilities of individual sensory nodes to concurrently execute multiple application tasks, allowing applications to group such sensory nodes together according to the requirements.

Our claim for the WSAN virtualization differs from the first three works discussed,[1–3] which are centralized virtualization models based on the traditional sensing-as-a-service model. Because our virtualization model is partly decentralized (based on the WSAN-as-a-service, or WSANaaS, model), this allows the execution of both centralized and decentralized applications and lower response time. Our proposed WSAN virtualization model also differs from Khan and his colleagues' approach, which is a partly decentralized virtualization model because it builds on the concept of information fusion.[4] Information fusion lets us reduce the amount of data manipulated by applications, which in turn saves energy for the physical WSAN and prolongs the lifespan.

References

1. D. Phan et al., "Toward Sensor-Cloud Integration as a Service," *Proc. 8th Int'l Conf. Body Area Networks*, 2013, pp. 355–362.
2. S. Madria et al., "Sensor Cloud: a Cloud of Virtual Sensors," *IEEE Software*, vol. 31, no. 2, 2013, pp. 70–77.
3. K. Kim et al., "Agriculture Sensor-Cloud Infrastructure and Routing Protocol in the Physical Sensor Network Layer," *Int'l J. Distributed Sensor Networks*, vol. 2014, 2014, pp. 1–12, article 437535.
4. I. Khan et al., "Wireless Sensor Network Virtualization," to be published in *IEEE Network*, May/June 2015.

IEEE CLOUD COMPUTING
Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page
Qmags THE WORLD'S NEWSSTAND®

as computational entities capable of performing a set of information-fusion techniques. Therefore, it's possible to map the virtualization of physical sensors onto the three data abstraction levels of information fusion (measurement, feature, and decision),[8] based on the input and output data of each virtual sensor.

The second major drawback is that such a communication overhead imposes an additional delay on top of the response time of applications running within the CoS. Application response time consists of the execution time of data acquisition, processing, decision, and actuation. In centralized CoS infrastructures, an application's response time depends on several factors,[2,4,9] such as

- the formation of communication bottlenecks close to the sink nodes,
- the size (in hops) of the physical WSAN,
- the delay in routing data inside the cloud,
- the delay in processing and making decisions within the cloud, and
- the delay in issuing an actuation message back to the physical WSAN.

Although the data reduction approach helps lessen the time spent due to each of these factors, when transmitting the data via a sink node to the cloud for instantiating virtual nodes and data processing, the total response time still comprises all such factors. Such a situation impedes several applications that require fast response.[10] To overcome this restriction, a feasible approach is to decentralize applications' decision processes—that is, perform application decision processes inside the physical sensor, leveraging the nodes' in-network processing capabilities.[6]

Furthermore, the existing publish/subscribe models proposed within the centralized CoS infrastructures ignore the fact that sensors also have local processing and communication capabilities for performing localized and collaborative algorithms, which are required for applications that are inherently decentralized.[6] In particular, the adoption of the WSAN as a cornerstone of the IoT paradigm fosters the introduction of novel and more complex applications, such as domotics, assisted living, e-health, business/process management, structural health monitoring, and intelligent transportation of people and goods. To complete complex tasks in the IoT scenario, applications require distributed processing within the network. In general, the centralized CoS infrastructure approach is unsuitable for the execution of decentralized applications. Therefore, to support a broader set of applications, the CoS infrastructure must allow the execution of localized and collaborative algorithms as a service within the physical sensors. Such an approach leads to a WSAN-as-a-service (WSANaaS) paradigm, in which the concept behind the services provided by a WSAN node is much broader than the concept proposed in the traditional sensing-as-a-service paradigm.[4]

Olympus is an information fusion and CoS-based decentralized WSAN virtualization model that seeks to make the best use of the cloud and the physical WSAN environments by finding a balance between two possible approaches for running services: centrally, inside the cloud, and locally, within the physical sensors. Olympus uses information fusion to ensure that the system will provide data at a given abstraction level of the manipulated data.[8] It's a decentralized WSAN virtualization model because physical nodes can perform the necessary procedures for creating and running the virtual sensor locally. Therefore, in Olympus, application decision processes are performed partly within physical sensors and partly within the cloud.

Information Fusion

According to Eduardo Nakamura and his colleagues, in the WSAN field, information fusion techniques are used to either reduce the communication overhead to reduce nodes' energy consumption or to improve the performance (accuracy) of the applications (information).[8] Here, accuracy can be defined as the degree of proximity between the observed measurement and the expected value.

One aspect we can use to categorize information fusion is the data abstraction level.[8] According to this categorization, information fusion can be classified into three levels:

- The *measurement* level deals with one- or multidimensional signals originating from the sensors (generally raw data). Raw data are provided as input to the information fusion process and combined with a new collection of more accurate data, possibly with lower noise.
- The *feature* level handles characteristics (attributes or features) that are extracted from signals.
- The *decision* level deals with the decisions or symbolic representations taken as inputs for making a more global or confident decision on the data samples.

Here, we use the data-feature-decision (DFD) model, which provides a classification for information fusion according to the abstraction of the input and output data.[8] In *data in-data out* (DAI-DAO), information fusion deals with measurement-level data as input and output. In *data in-feature out* (DAI-FEO), information fusion uses data at the measurement level as input to extract attributes or characteristics that describe more summarized information for a given monitored area. In the *feature in-feature out* (FEI-FEO) category, information fusion works on a set of features to improve or refine an existing characteristic or attribute, or to extract new ones. In the *feature in-decision out* (FEI-DEO) category, information fusion uses a number of features extracted for generating a symbolic representation (or a decision). In the *decision in-decision out* (DEI-DEO) category, decisions can be merged to obtain new decisions. Finally, in the *data in-decision out* (DAI-DEO), either a decision is made directly over raw data, as an atomic procedure, or the information fusion process under this category can be broken into several atomic parts pertaining to other categories.

Olympus is based on the concept of WSAN virtualization to support multiple sensing applications running within the CoS.

Olympus Framework

Olympus is based on the concept of WSAN virtualization to support multiple sensing applications running within the CoS. Here, we consider an application as a set of services that must be performed to accomplish the application goals. Each application considers a geographical area of interest and is considered to have a finite lifespan. An application defines a set of quality-of-service (QoS) requirements, described in terms of maximum end-to-end delay, maximum percentage of packet loss, and energy consumption.[11] Moreover, applications require a set of services provided by the physical WSAN nodes that are described in terms of the following provided services: data collection, processing, decision, routing, and actuation. Such capabilities must be published within the cloud in a central repository through a publish/subscribe mechanism.[2] However, physical sensors connected to the CoS must have the minimal capability of locally (in its physical location) providing continuous raw data at a periodic rate. This premise is less restrictive than the works proposing centralized CoS infrastructure support, in which the physical nodes must provide such raw data directly to the sink node.

To connect applications to physical WSAN nodes, one or more virtual WSANs must be created. A virtual WSAN is formed by providing logical connectivity among the physical nodes. Such physical nodes are grouped into different virtual WSANs based on the phenomenon being monitored or the service being provided.[5] Imran Khan and his colleagues identify two categories of WSAN virtualization: node level and network level.[12] Node-level virtualization allows multiple applications to run services concurrently on a single WSAN node. In network-level virtualization, a subset of physical WSAN nodes forms a virtual WSAN, made of virtual WSAN nodes, to execute given application tasks at a given time.

A virtual WSAN node is an abstraction of a set of physical nodes from which the virtual node obtains data. Such a virtual node is considered a computational entity capable of performing a set of information fusion techniques at a given level of the DFD model. Virtual nodes can also form logical neighborhoods. In contrast to physical neighborhoods, usually defined in terms of radio ranges, the nodes in a logical neighborhood are specified by the application based on specific requirements.[5] Our model includes a computational entity, the *virtualization manager*. This entity runs within the cloud and has several responsibilities regarding model execution management. Our model is said to be partly decentralized because the services allocated by this entity will run within the physical nodes.

Virtualization Based on Information Fusion

To build the proposed mapping, it's first necessary to abstract the physical world. That is, to abstract issues regarding the spatial/geographical distribution of each sensor, as well as to abstract which of the sensors pertain to each physical WSAN deployed in different locations. Consequently, every physical sensor (green circles in Figure 1) can be on the same physical WSAN as others or in separate physical WSANs. All the nodes comprise the physical layer, independently of their physical organization. In this physical layer, we consider a scenario with multiple sensor infrastructure providers.[5] Over the physical layer is the information fusion layer, which is divided into the measurement, feature, and decision levels, as discussed earlier. Now, we can represent all of the virtual sensors that can be logically formed by instantiation from a physical WSAN to fit each of the information fusion levels.

As Figure 1 shows, the virtual sensors depend not only on the information fusion level of the input/output data, but also on the source of this data within the physical network. According to the data source, we can define the following types of virtualization, inspired by Sanjay Madria and his colleagues: one-to-one, one-to-many, and many-to-one.[4] In one-to-one virtualization, the virtual sensor is created just to replicate a physical sensor's raw data (in the case of DAI-DAO virtualization). In the one-to-one case, it's also possible that an information fusion technique, such as filters or signal processing, is performed within this representation, which is also the case for DAI-DAO virtualization. In the one-to-many virtualization case, the many virtual sensors are created as different representations of the same original data. Such different representations can be achieved by performing different information fusion techniques over the same raw data within each of the virtual sensors. In many-to-one virtualization, in the case of the DAI-DAO virtualization, each virtual sensor could represent the arithmetic mean of two or more raw datasets from different physical sensors but obtained at the same time. This helps achieve redundancy and improve data accuracy.

The one-to-one, one-to-many, and many-to-one virtualizations can also be defined at the DAI-FEO level. The fundamental difference is that instead of providing the measurement data as output, the output data will be features, such as when a node calculates an arithmetic mean from a set of raw data obtained from one or more nodes.

All the types of virtualization described up to this point result from a primary virtualization process—that is, a virtualization performed when a first virtual sensor is instantiated directly from a physical sensor. The virtual sensors resulting from the primary virtualization processes are shown as red circles in Figure 1. From this point on, all the virtualization processes will consist of virtual sensors being created from other virtual sensors. Composed virtualizations can also occur as one-to-one, one-to-many, or many-to-one virtualizations. Such virtualizations can occur within the DAI-DAO and DAI-FEO information fusion levels as well, as the primary virtualization. However, only the composed virtualization processes can achieve higher levels, such as the FEI-FEO, FEI-DEO, and DEI-DEO.

It's important to mention that within the DAI-DAO, FEI-FEO, and DEI-DEO levels, several

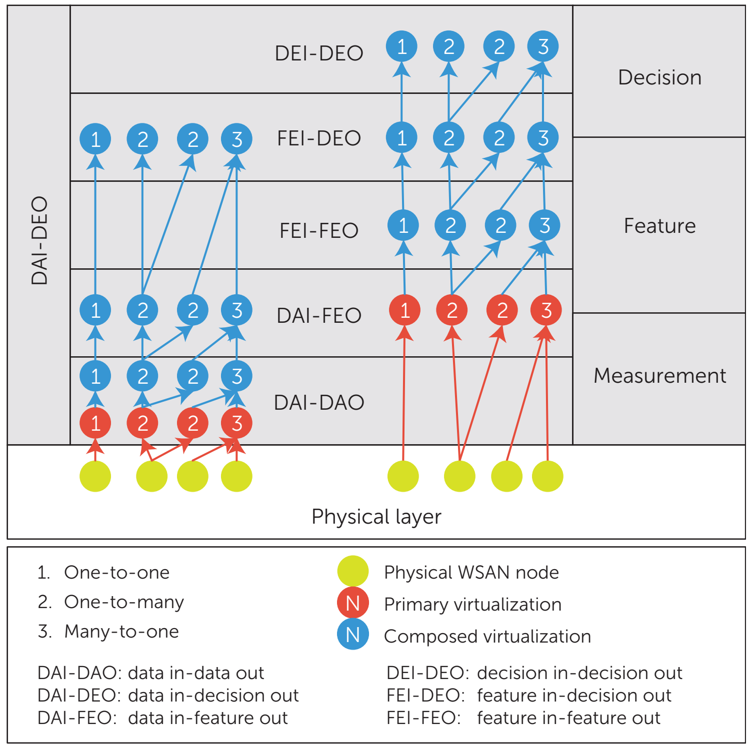

FIGURE 1. Linking virtual nodes and data abstraction levels of information fusion. The green circles represent every physical sensor, and the red circles represent the virtual sensors resulting from the primary virtualization processes.

compositions of virtual sensors can occur without changing the information fusion level. But inside the DAI-FEO, FEI-DEO, and DAI-DEO levels, if a composed virtualization is performed, then a composed virtualization starting from the initial virtualization will inevitably result in a change of information fusion level. The virtual nodes originating from the composed virtualizations can never be created directly from a physical sensor. The composed virtualizations must always originate from a DAI-DAO or DAI-FEO information fusion level.

Operational Model

During virtual WSAN node creation, the communication overhead among the physical sensors and the cloud tends to be higher (but for a shorter period) than during virtual sensor operation management. The first issue to deal with in this phase is to manage a publish/subscribe mechanism to choose the proper physical sensors that meet the requirements of the applications requesting the creation of the virtual WSAN node.[2]

First, every physical sensor publishes its capabilities within the cloud using push communication.[7]

Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

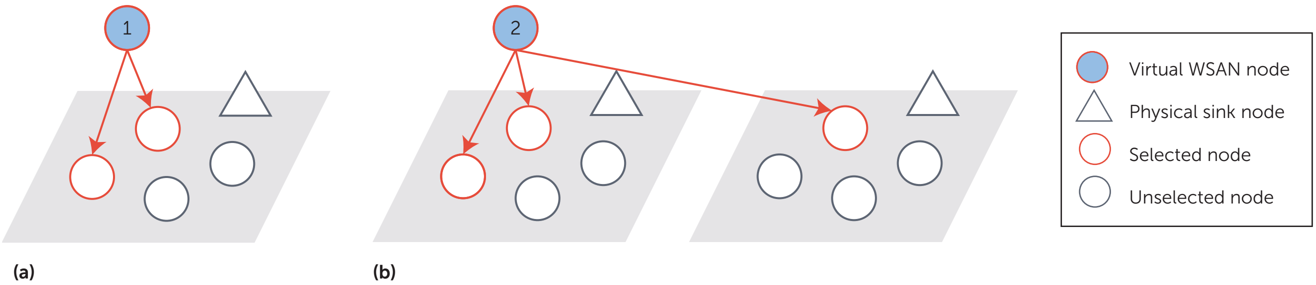

FIGURE 2. Examples of many-to-one virtualization (a) within the same physical wireless sensor and actuator network (WSAN) and (b) comprising different physical WSANs.

The virtualization manager is responsible for reading the application requirements (through subscriptions) and the published physical sensor capabilities, and deciding if a new instance of a virtual sensor should be created or an existing instance should be reused. The virtualization manager searches for the physical sensors fully inside the cloud (where the respective capabilities are published). However, when a new virtual WSAN node is instantiated, the virtualization manager must communicate with the respective physical sensor to establish routes and ensure that the published capabilities required by applications are still available. The virtualization manager also keeps track of the unique identification addresses of the virtual sensors and can download, to the physical sensors (through dynamic loading techniques), any software required to virtualize the physical node.

The physical sensors, on receiving any request (sent by the virtualization manager) to start the instantiation of a virtual node, are elected as leaders. In our model, a physical sensor is either a *leader* node or a *common* node. Leader nodes (elected by the virtualization manager) search for and establish routes for communicating with other physical sensors (possibly in other physical WSANs) required by a given virtual sensor. Leader nodes act as reference nodes for the virtualization manager to communicate with when retrieving data or performing procedures for virtual WSAN node creation and operation management. Leader nodes perform information fusion at the highest level required by the instantiated virtual node. Therefore, the output data of leader nodes is exactly the output data expected from the virtual sensors. Given the procedures described, we conclude that virtual sensor creation must be performed partly within the cloud and partly within the physical sensors.

When virtual WSAN node operation management starts, all physical sensors are ready to perform any service allocated to the respective virtual sensors. In this phase, the virtualization manager allocates execution of the application's services within the physical WSAN (service allocation).

Regarding the one-to-one and one-to-many virtualization types, we envision that such service allocation is relatively simple because all the physical nodes involved in the virtualization process are within the same physical WSAN. Therefore, all the communications are performed within the same WSAN, which in turn reduces the application response time because there's no need to route messages through the sink nodes and the cloud. Moreover, the only existing physical sensor is the elected leader node. For instance, one physical sensor can perform the processing needed for virtualization, collect and reduce data, and communicate the final features or decisions to the virtualization manager. If an undesired state is detected, the virtualization manager will issue warnings to the application users through push communication. However, every physical sensor must be ready to perform pull communication during virtual sensor operation management in case application users want to stay up to date with the execution status of the allocated services. When a virtual node is instantiated, the user watches its behavior within the cloud as an abstracted entity with updated status. This status must be updated according to the application requirements.

The many-to-one virtualization type, when consisting of nodes in the same physical WSAN (Figure 2a), is subjected to an analysis similar to that for the one-to-one and one-to-many virtualization types. The only difference is the higher communication overhead and delay for establishing routes among all the physical sensors that compose the virtual node. However, if any physical sensor pertains to a separate physical WSAN (Figure 2b), the routing among the physical sensors must go through the sink nodes of both physical WSANs to enable passage through the cloud. This situation is a critical issue. We consider two possibilities in our model to resolve it. The first approach is to treat the issue as a routing problem. The information must flow among the physical sensors of separate physical WSANs to reach the leader node representing the virtual node, which will perform the information fusion. In such a routing problem, the information will inevitably flow through two sink nodes (gateways); consequently, such routing must rely on the physical sensors as well as the cloud, with the virtualization manager acting as a coordinator. The second possibility is equivalent to the centralized CoS virtualization model approach. That is, the physical sensors send data to the cloud through the sink nodes, and virtual node creation takes place only within the cloud. The Olympus framework considers both possibilities. The virtualization manager must switch between them, depending on the current application requirements.

Finally, the virtual WSAN node operation is more prone to running within physical sensors than virtual sensor instantiation. Virtual sensors can even operate without any communication with the cloud. However, there are still virtual sensor operation procedures that can be performed partly within the cloud, such as the case of many-to-one virtualization comprising nodes physically separated in different WSANs. Thus, Olympus is considered a hybrid (partly decentralized) WSAN virtualization model.

Olympus extends the physical WSAN lifespan, reduces application response time, and supports several centralized and decentralized applications. Several directions deserve further study.

First, we need solutions for routing messages in Olympus, especially in the case of many-to-one virtualization with any physical sensor pertaining to a separate physical WSAN. Because in this case routes must pass through gateways (and so, through the cloud), it's worth considering routing solutions that involve the virtualization manager. Because the virtualization manager has access to all active applications in the CoS, application semantics (possibly by correlating data) might improve routing.

Another area of investigation is to identify an optimal point at which to apply data reduction (through information fusion) in the proposed virtualization model, relating the benefits achieved, in terms of a reduction in communication overhead, and the loss of accuracy.

We need solutions for forming logical neighborhoods among virtual WSAN nodes. In this context, we intend to formally describe a scheme for gathering information (rules), regarding logical neighborhood formation, from applications and for reflecting such information on the virtual WSAN. Solutions for the WSAN virtualization problem within the field of service allocation are also needed. In this sense, we'll formally describe a service allocation procedure to be performed by the WSAN virtualization manager in Olympus.

Finally, we plan to implement a system based on the Olympus model. We'll experiment by connecting it to WSAN testbeds in our university, and assess the WSAN lifespan and application response time achieved with our system in comparison to centralized CoS infrastructures.

Acknowledgments

This work was conducted during a scholarship supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) at Universidade Federal do Rio de Janeiro. This work was also supported in part by Conselho Nacional de Desensolvimento Científico e Tecnológico (CNPq) under grants 304941/2012-3, 473851/2012-1, 477223/2012-5, and 307378/2014-4; by Instituto Nacional de Metrologia, Qualidade e Tecnologia (Inmetro/Pronametro); and by Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ) under grant JCNE E-26/102.961/2012. Albert Zomaya's work is supported by Australian Research Council Discovery Grant DP130104591. Igor L. Santos is the corresponding author.

References

1. A. Botta et al., "On the Integration of Cloud Computing and Internet of Things," *Proc. Int'l Conf. Future IoT and Cloud* (FiCloud 14), 2014, pp. 23–30.
2. A. Alamri et al., "A Survey on Sensor-Cloud: Architecture, Applications, and Approaches," *Int'l J. Distributed Sensor Networks*, vol. 2013, 2013, article 917923.
3. L. Atzori et al., "The Internet of Things: A Survey," *Computer Networks*, vol. 54, no. 15, 2010, pp. 2787–2805.
4. S. Madria et al., "Sensor Cloud: A Cloud of Virtual Sensors," *IEEE Software*, vol. 31, no. 2, 2013, pp. 70–77.
5. M. Islam et al., "A Survey on Virtualization of Wireless Sensor Networks," *Sensors*, vol. 12, no. 2, 2012, pp. 2175–2207.
6. I. Dos Santos et al., "A Localized Algorithm for Structural Health Monitoring Using Wireless Sensor Networks," *Information Fusion*, vol. 15, Jan. 2014, pp. 114–129.

CLOUD COMPUTING Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

7. D. Phan et al., "Toward Sensor-Cloud Integration as a Service," *Proc. 8th Int'l Conf. Body Area Networks*, 2013, pp. 355–362.
8. E. Nakamura et al., "Information Fusion for Wireless Sensor Networks," *ACM Computing Surveys*, vol. 39, no. 3, 2007, pp. 1–55.
9. K. Kim et al., "Agriculture Sensor-Cloud Infrastructure and Routing Protocol in the Physical Sensor Network Layer," *Int'l J. Distributed Sensor Networks*, vol. 2014, 2014, pp. 1–12, article 437535.
10. S. Ajmal et al., "An Efficient Admission Control Algorithm for Virtual Sensor Networks," *Proc. IEEE Int'l Conf. High Performance Computing and Comm.*, 2014, pp. 735–742.
11. W. Li et al., "Efficient Allocation of Resources in Multiple Heterogeneous Wireless Sensor Networks," *J. Parallel and Distributed Computing*, vol. 74, no. 1, 1974, pp. 1775–1788.
12. I. Khan et al., "Wireless Sensor Network Virtualization: Early Architecture and Research Perspectives," to be published in *IEEE Network*, May/June 2015.

IGOR L. SANTOS *is a doctoral student in informatics at the Universidade Federal do Rio de Janeiro (UFRJ), supported through a scholarship by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), and a professor in the Computer Science Department at UFRJ. His research interests include wireless sensor networks, cloud computing, information fusion, and structural health monitoring. Santos has an MS in informatics from UFRJ. Contact him at igorlsantos@gmail.com.*

LUCI PIRMEZ *is a professor of postgraduation courses in computer science at the Universidade Federal do Rio de Janeiro (UFRJ). Her research interests include wireless networks, wireless sensor networks, network management, and security. Pirmez has a PhD in computer science from UFRJ. Contact her at luci.pirmez@gmail.com.*

FLAVIA C. DELICATO *is an associate professor of computer science at the Universidade Federal do Rio de Janeiro (UFRJ). Her research interests include wireless sensor networks, adaptive systems, middleware, and software engineering techniques applied to ubiquitous systems. Delicato has a PhD in electrical and computer engineering from UFRJ. She is a level 1 Researcher Fellow of the Brazilian National Council for Scientific and Technological Development. Contact her at fdelicato@gmail.com.*

SAMEE U. KHAN *is an associate professor of electrical and computer engineering at North Dakota State University. His research interests include optimization, robustness, and security of cloud, grid, cluster, and big data computing, social networks, wired and wireless networks, power systems, smart grids, and optical networks. Khan has a PhD in computer science from the University of Texas, Arlington. He's a senior member of IEEE. Contact him at samee.khan@ndsu.edu.*

ALBERT Y. ZOMAYA *is the chair professor of high-performance computing and networking in the School of Information Technologies at Sydney University. His research interests are in the areas of complex systems, parallel and distributed computing, and green computing. Zomaya is editor in chief of* IEEE Transactions on Computers *and is the recipient of the IEEE TCPP Outstanding Service Award and the IEEE TCSC Medal for Excellence in Scalable Computing. He has a PhD in control engineering from Sheffield University, UK. Zomaya is a Fellow of AAAS, IEEE, and IET (UK). Contact him at zomaya@it.usyd.edu.au.*