

Towards the Future Internet of Sensors

Olivier Alphand, Andrzej Duda, Martin Heusse, Benoît Ponsard, Franck
Rousseau, Fabrice Theoleyre

► **To cite this version:**

Olivier Alphand, Andrzej Duda, Martin Heusse, Benoît Ponsard, Franck Rousseau, et al.. Towards the Future Internet of Sensors. The Internet of Things, Springer New York, pp.309-318, 2009, 10.1007/978-1-4419-1674-7_30 . hal-01073336

HAL Id: hal-01073336

<https://hal.inria.fr/hal-01073336>

Submitted on 2 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards the Future Internet of Sensors

Olivier Alphand, Andrzej Duda, Martin Heusse, Benoît Ponsard, Franck Rousseau,
and Fabrice Theoleyre

Abstract In this position paper, we propose a new view on the integration of wireless *Sensor and Actuator Networks* (SANET) in the Internet. We believe that SANETs may benefit from the *data dissemination* paradigm in which the network conveys typed data chunks while nodes organize communication according to the Publish/Subscribe model: data consumers subscribe to chunks advertized by producers. A *data router* interconnects SANETs and offers a data centric view on the physical world to the rest of the Internet. Our paper presents the details of the data centric view on integrating SANETs into the Internet, discusses several related issues, and proposes a research agenda for future activities.

Key words: Wireless Sensor and Actuator Networks, content-centric networking, publish/subscribe model

1 Introduction

In this position paper, we propose a new view on the integration of wireless *Sensor and Actuator Networks* (SANET) in the Internet. Such networks become increasingly important for gathering various physical measures and acting upon objects in the physical world. Usually, they support a single application and use sophisticated protocols and wireless technologies optimized for energy consumption. Significant recent research effort aims at making them more generic and interconnecting with end-hosts on the Internet.

One approach to achieve this objective is to provide end-to-end IP connectivity to all sensor/actuator nodes. IETF promotes the idea of using the standard Inter-

Olivier Alphand, Andrzej Duda, Martin Heusse, Benoît Ponsard, Franck Rousseau, Fabrice Theoleyre
Grenoble Informatics Laboratory, University of Grenoble, Grenoble, France. e-mail: {firstname.lastname}@imag.fr

net protocols: the 6LowPAN working group proposes to adapt IPv6 to operate in a sensor network [8]. As the IPv6 addresses are long, the group proposes to set up a gateway that translates 128 bit addresses into 16 bits. Moreover, header compression techniques are further used to reduce the overhead. In addition to these techniques, the IETF ROLL working group proposes to develop a routing protocol over low power and lossy networks [6]. Several industrial developments (ArchRock, Dust Networks, Cisco) and alliances (IPSO Alliance - IP for Smart Objects) follow the same track. However, we think that it is also interesting to explore other approaches. As SANET nodes are very different from IP routers, we would like to find another paradigm that provides a higher-level, data oriented view on the physical world in which this kind of networks operate.

At the same time we observe that such a data centric view perfectly corresponds to the new approach of the Future Global Internet: several initiatives proposed to base the new generation of the communication infrastructure on the *data or content dissemination* paradigm. We believe that such a model is particularly suitable for SANETs and allows us to begin experimentation with the new vision of the Future Internet in the context not so ossified as the current Internet. At the same time, it may help gaining insight into new solutions to specific problems of SANETs. Our paper presents the initial ideas on the data centric view for integrating SANETs into the Internet, discusses several related issues, and proposes a research agenda for future activities.

The rest of the paper is organized as follows. We discuss SANETs and their integration in the current Internet in Section 2. Then, we present our view of the data centric approach to the future Internet of Sensors in Section 3. Section 4 briefly discusses more detailed issues of a research agenda for achieving this objective. Section 5 summarizes the related work and Section 6 concludes the paper.

2 SANET Networks and the Internet

Sensors and actuators are increasingly used in many applications for gathering information about the physical world or controlling it. We can cite industrial applications in which sensors measure operation parameters for detecting abnormal conditions and controlling affected systems. Smart buildings use sensors for controlling temperature, light intensity, humidity, or detecting a fire. Multimedia sensors are able to capture images or video streams for intrusion or feature detection, tracking objects, etc. Sensor networks communicate by means of cheap radio, which significantly reduces the deployment cost. In SANETs, some nodes have also the ability to act on their environments, for instance they can move objects as well as activate or control some devices. Sensor nodes face many challenges in networking, embedded systems, databases, and hardware design because they are highly constrained: their processing power is limited, they have small memory, and they are battery operated with the requirement of a long life. Their internal design, protocols at network and MAC layers, as well as application operation need to satisfy all these stringent con-

straints and provide a suitable tradeoff between small imprint, efficiency, and energy consumption.

Sensor networks have become a challenging hot research topic and one aspect of the current research is their integration with the Internet. IETF adopted a conservative view of applying the standard Internet protocols to SANETs thus providing the end-to-end IP connectivity to all sensor nodes. Pushing IPv6 to sensor nodes seen as end systems raises several issues. First, using the unchanged IPv6 results in too much overhead, so the version for sensor networks needs to use a shortened header as mentioned above and this modification means that there is no pure end-to-end connectivity, because packets need to be translated by a gateway acting as a sort of a NAT (Network Address Translation) box. Second, short 16 bits addresses may be not sufficient for large-scale sensor networks (if sensor nodes form a "dust", thousands of nodes need to be addressed). Furthermore, perhaps there is no need for assigning an IP address to each sensor node, because better views on the information provided by sensors are possible. The analogy that we take to explain this issue is a computer composed of a CPU, memory, network card, an external disk, and a graphic card. Why do not we assign an IP address to the graphic card? It contains a processor, some memory and we could provide IP connectivity to each computer component. Simply, we do not need such a feature, because the view provided by the computer to other end-systems is different and does not include communicating with a graphic card. The internal bus interconnects these components, they have various functions, and you do not need to extend IP connectivity to them. At the same time, the computer provides kind of a virtual view on all its components and it is viewed by other end-systems connected to the Internet as a unique entity.

Sensor networks are somehow similar—they need to provide the information gathered by sensors, but they can do it independently of considering each sensor node as an IP end-system. Furthermore, the stringent constraints of sensor nodes call for optimized MAC and routing protocols, as well as a cross layer approach for efficient interaction of their operation. The ROLL working group that proposes a "Routing Over Low power and Lossy networks protocol" has begun to address this issue [6]. In fact, SANETs are intrinsically different from the global Internet, because they operate according to a different paradigm. SANET nodes sleep almost all the time for saving energy and wake up to transmit low volumes of data, while the Internet supports high bandwidth planetary communication infrastructure. Imagine that many Internet hosts want to communicate with a single sensor node, this may lead to the performance bottleneck in the SANET. The situation of SANETs is similar to residential networks in which people do not run Web servers, but rather upload content to Web servers in ISP networks—content producers are not the same as content serving hosts. A special gateway between a SANET and the Internet may provide an abstract view on data gathered by sensor nodes, serve them for many destinations, and eliminate possible performance bottlenecks.

Note also that the sensor network community has already considered this architectural issue and converged to the conclusion that SANETs need to be seen as a Layer 2. Polastre et al. state that "*wireless sensor networks would benefit from a*

unifying abstraction (or "narrow waist" in architectural terms), and that this abstraction should be closer to the link level than the network level" [10].

3 Future Internet of Sensors

We start by considering a new vision of a communication substrate that a SANET network needs to provide to Internet end-systems. We can adopt a similar approach as the Internet at its beginning that aimed at providing a minimal scalable interconnection layer with a horizontal interface over different transmission technologies and enabling new communication applications. The socket layer to TCP/IP contributed to the initial success of the Internet by offering the right interface to applications. We thus would like to provide similar features in SANETs—interconnect them on a large scale and provide applications in the Internet with the ability to interact with the physical world.

We think that SANETs are very different from traditional end-points in the current Internet, because they provide the interface to the physical world. In the Internet, IP addresses belong to a structured virtual space unrelated to the physical world, while sensor nodes are placed at a given geographical location. For most applications the information on the place and the time of a sensor reading is almost as important as the reading itself. For instance, it would be important to know that a node has detected a fire in a given office at a given time. At the same time, much information gathered by sensors is asynchronous, for instance a node that detects a threshold crossing of a measured quantity needs to notify a base station for further processing. All these examples lead to the conclusion that the focus of communication in SANETs is on data, which can include other information such as place and time in addition to a simple reading, rather than on IP addressable end-points, so the current communication model of the Internet is not suitable for SANETs. Thus, we can imagine a layer that offers the view of typed data chunks coming from and going to some nodes instead of the traditional vision of addressable end-points. Briscoe has expressed a similar view on the interest of a data centric approach applied to sensor networks and pervasive computing [2].

Similarly, one would be interested in communicating with a SANET network through primitives that allow expressing the fact that a given piece of information comes from a given place, at a given instant, and provides a measurement of a given quantity. Current communication primitives such as sockets do not support this data-centric view.

To define the right place of SANETs in the Internet, we consider below the following most relevant issues:

- Interconnection architecture: what is a suitable architecture that provides minimal scalable interconnection?
- Communication interface: what is the right abstraction for representing a SANET network?
- Protocols: what are the right protocols to support the communication interface?

3.1 Interconnection Architecture

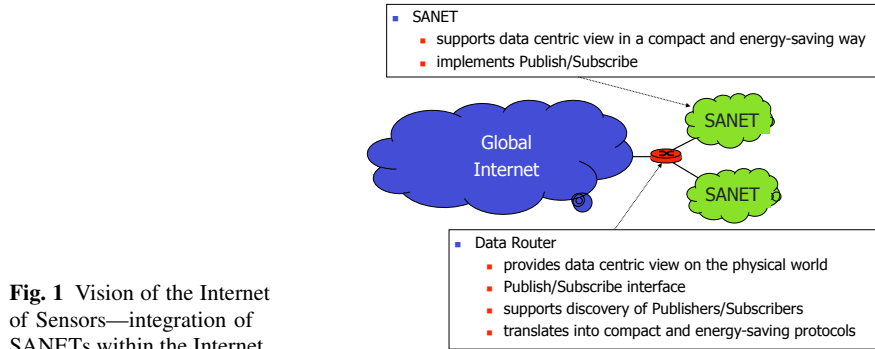


Fig. 1 Vision of the Internet of Sensors—integration of SANETs within the Internet.

Figure 1 illustrates the proposed interconnection architecture: a *data router* interconnects SANETs and offers a data centric view on the physical world to the rest of the Internet. As in the data centric view for the Future Internet, the data router operates on *data chunks* that are *typed*, *signed*, and possibly *encrypted*. Instead of communication end-points with IP addresses, *data publishers* representing sensor sources of information provide data chunks to *data subscribers*. Actuators can correspond to data subscribers willing to receive some data of special type, e.g. control commands. The data router acts as a data forwarder and a cache so that a subscriber can obtain a data chunk published before a subscription. It can interconnect several heterogeneous SANETs as well as offer data chunk publishing and subscription mechanisms to end-hosts in the Internet. To discover potential data publishers and subscribers, the data router offers a function that advertizes data types and metadata related to each publisher or subscriber. In this way an application can discover how a given data chunk can be consumed when received or what data chunk can be sent




Fig. 2 Data publishers and subscribers.

to a subscriber. The data router can also provide a directory service for discovering potential data publishers and subscribers. Note that the main difference between this view and existing Publish/Subscribe proposals for sensor networks [13, 4, 7, 9] is the placement of the functionalities in the protocol stack: in our vision, Publish/Subscribe protocols belong to lower layers and they are closely coupled with MAC access methods rather than provided by a middleware.

Figure 2 further illustrates the operation of the proposed interconnection architecture. We distinguish four entities: Sensors, Actuators, Publishers, and Subscribers (note that Sensor are also Publishers and Actuators are also Subscribers). Sensors publish their data for consumption by Subscribers and Publishers provide some data such as commands to Actuators. In the figure, we have placed Publishers and Subscribers in the Internet part, but they can also connect to a SANET. There can be several Data Routers and SANETs.

To explain the operation of the architecture, let us consider a sensor that publishes a data chunk: it propagates in the SANET (at the beginning it may propagate as a flooding message across the SANET) and reaches the Data Router that can store it in a cache and register in a directory of published content. The Data Router replies with a message that follows the inverse path to the Sensor thus fixing the direct route in intermediate nodes (note that this way of operation is inverse to Directed Diffusion [3]). A Subscriber sends request for a data chunk to the Data Router that can provide the data. If the operation starts with a Subscriber, it is up to the Data Router to propagate the request for data in the SANET so that the Sensor can publish its data chunks (this behavior is similar to Directed Diffusion). In a similar way, an Actuator can send a subscription for commands that the Data Router can match to the data provided by a Publisher.

Propagation of subscriptions in the Internet may rely on the standard IP multicast whereas routing in the SANET requires exploring new approaches, for instance based on establishing routing indices in intermediate nodes in function of received subscriptions and published data chunks. Routes need to take into account quality of radio links and energy of nodes. Propagation of subscriptions can also be related with setting up common time schedules for wake up at MAC layer. As we can imagine Data Routers that change their positions, routing in the SANET needs to dynamically adapt to their possible different locations.

3.2 DATA Layer and its interface

Figure 3 shows the structuring of the protocol layers. The *DATA* layer implements the Publish/Subscribe protocol for discovering and matching producers and consumers of typed data chunks. It offers the interface that we call *content* to emphasize on its role—enable communication via typed data content. Above the *DATA* layer, application protocols may provide some other useful high-level functions such as a database interface for *querying* information retrieved from a SANET network,

receiving *asynchronous notifications* or *pushing* some information to a SANET network. We can observe that Demmer et al. have defined a similar communication API [5], however it is more oriented towards the standard networks as a kind of a data-centric replacement to the socket layer.

Data chunks in our model are vectors of typed values: $\vec{V} : \{v_1, v_2, \dots, v_n\}$, where each component v_i is a value of a given type. For instance a data chunk published by sensor measuring a temperature at a given place at a given instant could be the following: $\{\text{Temp: int, Place: geo-coordinates, Time: timestamp}\}$; an actuator that controls a camera can be seen as a subscriber that can accept the following chunks $\{\text{Tilt: tilt-units, Pan: pan-units, Zoom: zoom-units, Place: geo-coordinates, Time: timestamp}\}$.

A flow of data chunks is uniquely identified by a Content Identifier (CID), a short value derived either from the identifier of a sensor (e.g. public key) or from some values of the data chunks that remains constant. The idea is to attach a small identifier to data chunks that the network could use for forwarding and dissemination. Depending on different types of data, CID can support different communication primitives:

- *geocast*: CID specifies geographic coordinates (e.g. camera at Tour Eiffel),
- *timecast*: CID specifies a given time instant (e.g. sensor data generated at midnight),
- *subjectcast*: CID specifies a name of an object or class of devices (e.g. a light switch)
- *controlcast*: CID specifies an operation to perform (e.g. camera tilt)
- *objectcast*: CID specifies the object of an operation (e.g. zoom on a car)
- *modecast*: CID specifies the mode of performing an operation. (e.g. wide pan)

The Data Router maintains the mapping between a CID and typed values of data chunks \vec{V} so that subscribers can discover the format of the information and learn how to interpret them. Another possibility to explore is to build a distributed service based on a DHT to store the mapping.

We can think of CIDs as a kind of generalized label used for forwarding data chunks. Consider the example of CID referring to geographical coordinates—they allow to exploit geographical routing that presents many advantages with respect to traditional network addressing. For instance, we can apply the approach based on Waypoint Routing [12] to SANETs. Similarly, we can explore how to set up




Fig. 3 DATA layer in the SANET protocol stack.

efficient and energy aware routing based on other types of information such as device names or operations to perform. Note that CIDs are intended to be only used within a SANET so that their scope is limited and particular routing protocols can be adopted [1].

To be able to subscribe to a given flow of data chunks, subscribers need to first discover them. Data Routers can run a Discovery service as a centralized directory or as a distributed service based on a DHT approach. Matching types advertized by Data Routers may require the use of *ontologies* to set formal basis on the terms used in type definitions and avoid problems with not exactly the same meaning of data chunks.

The main functions of the DATA layer protocol are twofold:

1. A Data Router needs to advertize flows of data chunks, their data types, and metadata describing the contents. Consider for example a sensor that measures a temperature at a given place and an actuator that adjusts heating to obtain a given temperature. A control application needs to discover the possibility offered by the sensor of providing the temperature reading and by the actuator of accepting to set a heating level.
2. Publishers and subscribers need to declare themselves through a Publish/Subscribe mechanism. For instance, a control application may subscribe to the data published by a sensor. Similarly, an actuator may need to subscribe to a possible source of heating control commands and a control application publishes such commands.

4 Research agenda

We point out here some issues of a research agenda to achieve the objectives of the Internet of Sensors.

- Design and develop a prototype Data Router for interconnecting SANETs and representing them in the Internet. As most of SANETs we consider are wireless, this means that the router is in fact a radio gateway between different technologies used in sensor networks (e.g. 802.15.4, Coronis Wavenis) and 802.11 WLAN connected to the Internet. We plan to develop such a prototype based on a software/reconfigurable radio able to operate as a 802.11 and a 802.15.4 entity. The card will use a FPGA software radio prototype and enable experiments with interfacing SANETs with the Internet through a 802.11 access network.
- Prototype the DATA layer for interconnecting with Internet applications. This will require an efficient implementation of the Publish/Subscribe protocols in the Internet part of the whole interconnection architecture. A natural support for his kind of protocols is IP multicast. The data router would need to provide support for advertizing flows of data chunks, informing about the types of their data, and supporting the data-oriented communication primitives discussed above such as geocast, timecast etc.

- Define and develop the internal routing and MAC protocols for supporting the data centric view inside SANETs. We plan to work on efficient and energy limited mechanisms for the Publish/Subscribe protocols based on CID forwarding. In particular, we think that forwarding published data chunks and subscribe messages may require setting up routing indices that dynamically provide the information about the neighbor to which a given message needs to go in function of its CID. Propagation of data chunks will also require efficient dissemination protocols such as probabilistic flooding or other optimizations that take into account energy. The internal SANET protocols would also require exploring new cross-layer approaches (PHY/MAC/NET) for finding the right trade-off between high performance and energy savings in the implementation of the discovery protocol, subscription propagation, and publishing data chunks. Supporting data-oriented communication primitives like geocast, timecast etc. inside SANETs may require providing geographical routing, localization schemes, time synchronization, and efficient packet forwarding.
- Specify security mechanisms for authenticating all entities, digitally signing and encrypting, if needed, data chunks, managing access authorizations in the Discover/Publish/Subscribe protocols.

5 Related Work

Several authors have considered the Publish/Subscribe model for sensor networks. Briscoe has discussed many issues related to this view in the context of ubiquitous networks and proposed several original ideas on how to integrate small sensor nodes within the large scale Internet [2]. Several authors have proposed to develop a Publish/Subscribe middleware for managing sensor networks [13, 4, 7, 9]. This approach integrates the Publish/Subscribe paradigm at upper layers so that it is difficult to achieve sufficient efficiency and energy aware operation in SANETs. RTFM (Rendezvous, Topology, Forwarding, and physical Media architecture) is the only architecture to our knowledge that puts the Publish/Subscribe into lower layers (on top of the link abstraction) [11], however this proposal concerns the Future Global Internet and not sensor networks.

6 Conclusion

In this position paper, we have proposed a new view on the integration of wireless *Sensor and Actuator Networks* in the Internet. We believe that SANETs may benefit from the *data dissemination* paradigm that some researchers suggest for the Future Internet. In this approach, the network conveys typed data chunks while applications organize communication according to the Publish/Subscribe model: data consumers subscribe to chunks advertised by producers. We think that such a model

is particularly suitable for representing the operation of SANETs and will foster new interesting research activities required for integrating them within the Internet.

In our view, the central element of the integrated architecture is a *data router* that interconnects SANETs and offers a data centric view on the physical world to the rest of the Internet. To discover potential data publishers and subscribers, the Data Router offers a function that advertizes data types and metadata related to each publishers or subscribers. At the same time, it needs to interact with sensor and actuators according to a lightweight Publish/Subscribe protocol tailored to the energy and computing constraints.

Acknowledgments

This work was partially supported by the French Ministry of Research project ARESA under contract ANR-05-RNRT-01703.

References

1. Bachir, A., Barthel, D.: Localized Max-min Remaining Energy Routing for WSN Using Delay Control. In: Proceedings of the ICC, IEEE International Conference on Communications, vol. 5, pp. 3302–3306 Vol. 5 (2005)
2. Briscoe, R.: The Implications of Pervasive Computing on Network Design. BT Technology Journal **22**(3), 170–190 (2004)
3. C. Intagonwiwat et al.: Directed Diffusion for Wireless Sensor Networking. IEEE/ACM Transaction on Networking **11**(1), 2–16 (2003)
4. Costa, P., Picco, G.P., Rossetto, S.: Publish-Subscribe on Sensor Networks: A Semi-probabilistic Approach. In: Proceedings of the 2nd IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS05). Washington DC, USA (2005)
5. Demmer, M., Fall, K., Koponen, T., Shenker, S.: Towards a Modern Communications API. In: Proceedings of HotNets-VI (2007)
6. Dohler, M., Watteyne, T., Winter, T.: Urban WSNs Routing Requirements in Low Power and Lossy Networks. Tech. rep., IETF ROLL workgroup (2008)
7. Hauer, J.H., Handziski, V., Köpke, A., Willig, A., Wolisz, A.: A Component Framework for Content-based Publish/Subscribe in Sensor Networks. In: Proc. of 5th European Workshop on Wireless Sensor Networks (EWSN). Bologna, Italy (2008)
8. Hui, J.W., Culler, D.E.: Extending IP to Low-Power, Wireless Personal Area Networks. Internet Computing, IEEE **12**(4), 37–45 (2008)
9. Nam, C.S., Jeong, H.J., Shin, D.R.: Design and Implementation of the Publish/Subscribe Middleware for Wireless Sensor Networks. International Conference on Networked Computing and Advanced Information Management **1**, 270–273 (2008)
10. Polastre, J., Hui, J., Levis, P., Zhao, J., Culler, D., Shenker, S., Stoica, I.: A Unifying Link Abstraction for Wireless Sensor Networks. In: ACM SenSys '05: Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems, pp. 76–89 (2005)
11. Sarela, M., Rinta-aho, T., Tarkoma, S.: RTFM: Publish/Subscribe Internetworking Architecture. In: Proceedings of the IST Mobile Summit. Stockholm, Sweden (2008)
12. Schiller, E., Starzetz, P., Rousseau, F., Duda, A.: Binary Waypoint Geographical Routing in Wireless Mesh Networks. In: Proceedings of the 11th International Workshop on Model-

- ing Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2008), pp. 252–259. ACM, Vancouver, Canada (2008)
13. Souto, E., Guimaraes, G., Vasconcelos, G., Vieira, M., Rosa, N., Ferraz, C.: A Message-Oriented Middleware for Sensor Networks. In: MPAC '04: Proceedings of the 2nd Workshop on Middleware for Pervasive and Ad-hoc Computing, pp. 127–134 (2004)