# Challenge: NOWiSeNets: A Network Of Wireless Sensor Networks for Internetworking the Physical World 

D.M. Blough ${ }^{*} \quad$ M. Leoncini ${ }^{\dagger} \quad$ G. Resta ${ }^{\ddagger} \quad$ P. Santi ${ }^{\ddagger}$


#### Abstract

Wireless Sensor Networks (WiSeNets), in which several tiny, wireless sensor nodes are connected in a multi-hop fashion have been subject of intensive research in recent years. In a near future, it is expected that several WiSeNets will actually be deployed for environmental monitoring, smart home environments, intrusion detection systems, factory monitoring, and so on, and the issue of how to deal with interactions between different WiSeNets will naturally arise. In this paper, we propose a vision of our future WiSeNet-instrumented world, in which WiSeNets (each of which composed of up to several thousands of sensor nodes) are networked together. This Network Of WiSeNets (NOWiSeNets) can be seen as extending the concept of internetworking to the physical world. The goal of this paper is to describe our view of the NOWiSeNets paradigm, to provide examples that show the potential advantages of interconnecting WiSeNets, and to discuss the many research challenges that must be solved to realize this paradigm.


## 1 Introduction

Wireless Sensor Networks (WiSeNets) have recently gained increased interest from the research community as well as industry. WiSeNets are formed by several "smart sensor nodes" (sensor node for short), which are connected through wireless, multi-hop connections. A sensor node is composed of one or more physical sensors (e.g. thermal, acoustic, pressure, and so on), a simple processor, limited-capacity memory, and a low-power radio. Thanks to the advances in miniaturization, a sensor node (including batteries) can be packed into a coin size board.
Sensor nodes in a WiSeNet exchange with each other the information collected from the environment through the physical sensors, process this information, and convey it towards one or more base stations (or gateway nodes), which provide the interface between the WiSeNet and the external user. Thanks to the information exchange/processing performed by the network

[^0]nodes, the user can obtain a very accurate picture of the phenomenon under consideration, considerably improving the observation accuracy achieved by traditional techniques (direct observation, fixed sensing stations, and so on).

The application scenarios for WiSeNets abound, ranging from smart home environments to intrusion detection systems, from large-scale environmental monitoring to tracking of wild animals in natural parks. In other words, WiSeNets have the potential of revolutionizing the way natural phenomena and human activities are monitored in the next few years. This explains the considerable interest of the research community in Wireless Sensor Networks.
WiSeNets are also receiving increasing interest from industry. Smart sensor nodes are being produced and commercialized by some electronic manufacturers. We cite Crossbow [8], a company that produces on a large scale the Motes sensor nodes developed at UC Berkeley. Other major silicon companies such as Intel, Philips, Siemens, STMicrolectronics, etc. are interested in the WiSeNet technology, and are developing their own smart sensor node platform.

There is also a considerable standardization activity in the field of WiSeNets. The most notable effort in this direction is the IEEE 802.15 .4 standard currently under development, which defines the physical and MAC layer protocols for remote monitoring and control, as well as sensor network applications. ZigBee [9] is an industry consortium (currently involving more than 100 members, representing 22 countries on 4 continents) with the goal of promoting the IEEE 802.15.4 standard.

Currently, we are in a phase in which the WiSeNet technology is rapidly maturing, but few applications based on sensor networks have been defined to date. In particular, industries strive to find significant markets for WiSeNets applications. The most promising ones seem to be home control, building automation, industrial automation, and automotive applications [9]. Nevertheless, the market for wireless sensor hardware is expected to grow at the rate of $20 \%$ per year in the next years, three times the growth rate of the wired sensor market [10].
Given the above discussion, it is easy to envision that in a not too far future thousands (millions?) of WiSeNets will be deployed worldwide, and the is-


Figure 1: Current view of tomorrow's world with WiSeNets.
sue of how to deal with interactions between different WiSeNets will naturally emerge. The easiest way to deal with this issue is to avoid and/or forbid the interaction between WiSeNets, motivated, for instance, by privacy and security reasons. However, it is our strong belief that there exists a better way to deal with the problem of WiSeNets interaction: instead of avoiding/forbidding interactions between WiSeNets, we propose to build a framework that allows a fruitful exploitation of these interactions, leading to the formation of a potentially much more powerful system: a Network Of WiSeNets (NOWiSeNets).

The goal of this paper is to describe our view of the NOWiSeNets paradigm, to provide examples that show the potential advantages of internetworking WiSeNets, and to discuss the many research challenges that must be solved to realize this paradigm.

## 2 Overview of the NOWiSeNets Paradigm

So far, research on WiSeNets has focused on challenges related with performing 'in network' tasks (routing of data packets within the network, data aggregation, query processing, etc.), and with providing an interface between the network and the external user, who is typically assumed to access the data generated by the WiSeNet through the Internet. The focus of current research activity on 'in network' tasks is due to the fact that many problems related to WiSeNets implementation are still to be solved.
Summarizing, the current view of our future world with WiSeNets is depicted in Figure 1: several WiSeNets, each of which acting independently of the others, are connected to the Internet through one or more gateway nodes. The user (provided she has the rights to access the data) can gather information from the sensor network by accessing the gateway node(s), either directly (e.g., through direct wireless communication), or remotely through the Internet. The typical example could be a WiSeNet used for smart home en-


Figure 2: Our view of tomorrow's world with WiSeNets: the NOWiSeNets paradigm.
vironments, which can be accessed by the owner either directly when she is at home, or remotely through the Internet when working, traveling, and so on. In this sense, WiSeNets are claimed to "extend the existing Internet deep into the physical world" [11]. However, in the current view of tomorrow's world, "extending Internet to the physical world" simply means that the Internet can be used to access data gathered from the physical world and/or control actuator devices remotely through the various attached WiSeNets.

The idea that motivates this paper is a different view of tomorrow's WiSeNets-instrumented world: instead of treating a WiSeNet as a stand-alone entity, which is interfaced with the external world through few gateway nodes, we envision a future in which WiSeNets are autonomous, (partially) open entities that communicate directly with each other, possibly without human intervention. Our vision of tomorrow's world is depicted in Figure 2: beside being possibly connected to the Internet, WiSeNets have the possibility of communicating directly with each other, possibly along multihop paths. Direct communication between WiSeNets can occur on a regular basis (for instance, between the smart home networks of nearby houses), or occasionally (for instance, when a seagull who is part of an animal tracking WiSeNet is within reach of one or more sensors of an ocean monitoring WiSeNet). We call this Network Of Wireless Sensor Networks NOWiSeNets. Note that in the NOWiSeNets model depicted in Figure 2, a WiSeNet might be (occasionally) disconnected from the Internet, and it might use other WiSeNets within reach to access the Internet through multi-hop paths. This feature might be particularly important when fast deployment of WiSeNets is required (for instance, in case of disaster events), and setting up an Internet connection to/from the deployed WiSeNet is costly and/or infeasible. In this scenario, using existing WiSeNets as


Figure 3: Communications between spatially overlapped WiSeNets take place through the Interconnecting Nodes (INs).
relay might ensure Internet connection with no need of additional infrastructure.

An example of interaction between different spatially overlapped WiSeNets is reported in Figure 3. We have a home sensor network, which has a gateway to the Internet. A few members of a bird movement tracking WiSeNet come into reach with some of the sensors of the home net, and communications between the two WiSeNets take place through one or more pairs of nodes that act as gateways between the two networks. To distinguish these nodes from the gateways to the Internet, we call them Interconnecting Nodes (INs). The bird sensor network can also interact with a sensor network used for prompt forest fire detection, through other INs. Thanks to the INs, the bird sensor network can deliver data to the Internet, although it does not have a direct communication with a gateway node.

## 3 Motivating scenario

As an example motivating our idea of interconnecting WiSeNets, we borrow a scenario described in [11]. A large-scale hazardous chemical gas leak occurs at a chemical processing plant near the city of Futuria. The authorities, promptly alerted by the WiSeNet that monitors the plant, decide to deploy an aerial WiSeNet to get a real-time situational assessment of the extent and movement of the gas release. The aerial WiSeNet is deployed by few unmanned airplanes which drop a few thousands of sensor nodes directly into the plume. The air-dropped sensors quickly self-organize into a WiSeNet, and start gathering data through chemical sensors. The collected data is conveyed towards the unmanned airplanes, which act as gateway nodes. So far, we have described a scenario which is coherent with the view of tomorrow's world summarized in Figure 1: the data acquired by the aerial WiSeNet is used by the local authorities to help planning the evacuation of Fu-
turia, and to evaluate the effects of the gas leak on the environment. Let us now make a step forward in depicting this scenario, assuming the various WiSeNets that are present in the environment can interact with each other, possibly without human intervention. As the plume moves, some of the sensor nodes of the aerial WiSeNet may come into reach of nodes belonging to different WiSeNets. For instance, with roof-top nodes of a smart home WiSeNet. As the nodes of the two networks are within each other radio range, they exchange information. In particular, the nodes of the aerial WiSeNet inform the nodes of the smart home WiSeNet of the gas leak, and of the potential danger for the humans and the environment. As a response to this alert, the smart home WiSeNet, through appropriate actuators, close the windows, shutters, and so on, and promptly inform the house owner of the dangerous situation. Since the owner is not at home, the alert is issued through the gateway node, which sends an SMS to her cell phone. Furthermore, the smart home WiSeNet informs the smart home WiSeNets of the nearby houses of the danger, so that they can undertake the adequate countermeasures even if they are not directly within reach of the aerial WiSeNet. Meanwhile, a seagull belonging to a wild animal tracking WiSeNet enters into the radio range of one or more nodes of the aerial WiSeNet, and exchange information about the weather conditions (wind intensity and direction, pressure, humidity and so on) that the plume will encounter a few kilometers away. This information is promptly forwarded to the gateway nodes of the aerial WiSeNet (the unmanned airplanes), which inform the authorities of the expected propagation path of the gas leak.

The scenario described above gives an idea of the potential of the NOWiSeNet paradigm: by dynamically interconnecting WiSeNets, and by promptly undertaking actions (possibly without human intervention), the
effects on humans and things of disaster events can be considerably mitigated. This would be only partially possible with the current view of tomorrow's WiSeNet instrumented world, in which the WiSeNet is seen as a stand-alone network whose primary purpose is to report to the external user observations about the monitored phenomenon.

It is easy to figure out many other scenarios in which WiSeNet internetworking proves very useful. We just briefly mention a few others.

- Sensor fusion. Two or more WiSeNets can be deployed in the same area (but at different times and/or by different people) to perform specific sensing tasks (e.g., acoustic and thermal monitoring). Appropriate combinations of the monitored events might help in reporting more accurate and/or complete information. While this could be possibly done by separately gathering the data (via the respective gateways), in-network information processing and aggregation can help in optimizing the communication resources, much as is currently prescribed for today's WiSeNets.
- Tracking a mobile agent in a possibly (much) wider area. For example, a burglar sneaking into an apartment can be detected by the local WiSeNets. Police, promptly informed, can take advantage of networked WiSeNets to better track the criminal's movements and catch him.
- Speeding up gas (or other chemicals) leaks alarms. Besides the detailed example made above, this is another possible application connected to smart home environments. Suppose a leak occurs in two nearby houses. If the respective WiSeNets are networked, prompt safety actions can be taken (e.g., switching the electricity off).


## 4 Research challenges

### 4.1 Definition of the NOWiSeNets paradigm

Several research challenges are related to the definition of the NOWiSeNets paradigm itself.

Defining the WiSeNet identity. First of all, the idea of 'interconnecting WiSeNets' implies that it must be carefully defined what constitutes a WiSeNet in our framework. In fact, if WiSeNets are considered as stand-alone networks (possibly interfaced with the Internet through one or more gateway nodes), as it is the case in current view of WiSeNets, an informal definition such as "a network composed of smart sensor nodes that wirelessly communicate with each other possibly along multi-hop paths" is sufficient to characterize the concept of WiSeNet. However, if we consider internetworking WiSeNets, we must provide a more detailed definition of what constitutes a WiSeNet. In particular, we must define the identity of a WiSeNet, which
allows to uniquely identifying that particular WiSeNet into the NOWiSeNets. Making an analogy with the Internet, we must provide an abstraction similar to the IP address used to uniquely identify a device in the Internet. A peculiarity of the NOWiSeNets paradigm with respect to the Internet is that each network component is a network itself, composed of up to several thousands of nodes.
The WiSeNet ID should be used to uniquely address each sensor network into the NOWiSeNets, but also to provide the sensor nodes of a certain WiSeNet with a way of recognizing each other as members belonging to the same WiSeNet. The issue of uniquely assigning sensor nodes to a WiSeNet is particularly relevant in case of spatially overlapped WiSeNets (see Figure 3): it is important that the sensor nodes are able to identify each other as belonging to the same WiSeNet, or to different WiSeNets. Note that the problem of dealing with spatially overlapped sensor networks will become more and more relevant in a near future, as the number of deployed WiSeNets will increase.

In principle, we might use the following method to manage WiSeNet IDs. Before deploying a WiSeNet, the owner (public organization, private company, private user, and so on) must obtain a unique and certified ID for the network. WiSeNet IDs are managed by a centralized authority, which maintains a WiSeNet registry and assigns certified IDs. Once the WiSeNet owner has obtained a certified ID for his network, she can embed this information into every sensor of her WiSeNet, so that the sensor nodes are able to recognize each other as belonging to the same network. At the same time, the WiSeNet ID uniquely identify the newly deployed network into NOWiSeNets.

Besides the ID, we might think that a set of attributes is associated with a certain WiSeNet. The attributes, which can dynamically evolve during the network lifetime, describe the features of the WiSeNet, and can be used to negotiate the access level to each other resources when different WiSeNets come into reach (see Section 4.4). In general, we can have static attributes, such as the network owner, the network functionality (e.g. home sensor network, or animal tracking sensor network, etc.), size, and so on; and dynamic attributes, such as those describing the connectivity of the WiSeNet. These network-level attributes might include an attribute that indicates whether the network has direct connectivity to the Internet (or the distance to the closest gateway to the Internet), and an attribute (or a list of attributes) describing the WiSeNets (i.e. their IDs) with which there exists a direct connection. Thus, a particular WiSeNet into NOWiSeNets is fully characterized by its ID and by the set of its attributes.

Designing the addressing mechanism. Another important research challenge is the design of the mech-
anism used to address entities in NOWiSeNets. Since NOWiSeNets entities are network themselves, a twolevel addressing mechanisms can be envisioned: at the higher level there is WiSeNet addressing, and at the lower level there is sensor node addressing.

The WiSeNet address is essentially its ID, which allows the unique identification of the network into NOWiSeNets. In case the WiSeNet is stationary (say, it is a smart home sensor net), its ID can be stored in some routing tables, so that it can be easily reached by other NOWiSeNets entities. The situation is more complex in case of a mobile WiSeNets (e.g. a flock of sensorequipped birds), since in this case dynamic addressing techniques must be implemented.

One possible solution to deal with this problem is to use a Mobile IP-like mechanism: there is a 'home agent' associated with a mobile WiSeNet, which is statically assigned with an address. The mobile WiSeNet, whenever possible (i.e. when there is connectivity), sends updates to the home agent regarding its position. When a user or another WiSeNet wants to send/retrieve data to/from the mobile WiSeNet, she/it contacts the correspondent home agent, which forwards the request to the last location where the mobile WiSeNet was detected.

Another possible solution to address mobile WiSeNets is appropriate for 'push' type of networks, i.e. for WiSeNets that regularly send data to the external user. In case of a mobile 'push' WiSeNet (e.g. a flock of sensor-equipped birds), we might think that the WiSeNet, whenever possible (i.e. when it is able to reach a gateway to the Internet) dumps data to, say, an Internet-connected machine. In this way, any user who wants to get the data collected by the mobile WiSeNet can simply access the correspondent Internet-connected machine.

Dealing with sensor node-level addressing is probably simpler, as in principle individual WiSeNets are free to use the sensor node addressing mechanism that best fits to the specific network functionality. In fact, in our view of the NOWiSeNets paradigm, entities which are external to a certain WiSeNet are not allowed (nor interested) to address a specific node in the network ${ }^{1}$. Thus, the specific mechanism used to address sensor nodes in a WiSeNet is transparent to the other NOWiSeNets entities.

Cooperation issues. Another challenge that arises in the definition of the NOWiSeNets paradigm is stimulating cooperation between WiSeNets. When sensor networks are considered as stand alone networks, it is natural to assume that the network nodes cooperate to perform a certain task: in fact, the sensor nodes belong to the same authority (the owner),

[^1]and cooperation is the natural node behavior. However, when several WiSeNets which are owned by different authorities interact with each other, it is expected that 'selfish' WiSeNet behaviors manifest themselves. In fact, a WiSeNet might have no interest in letting other WiSeNets using its (very limited) resources. A very similar problem arises in wireless ad hoc networks, where the network nodes are in general owned by different authorities. Several techniques to stimulate cooperation in ad hoc networks have been proposed in the literature (see, for instance, $[1,2]$ ), and we believe that they can be used as a starting point to devise similar mechanisms in the NOWiSeNets framework. However, we want to outline that the analogy between ad hoc networks and the NOWiSeNets paradigm (ad hoc network node $=$ WiSeNet) is somewhat imperfect: in fact, a WiSeNet is composed of many sensor nodes, which in general do not act as a unique entity as it is the case for an ad hoc network node. For instance, the decision about whether to relay a packet coming from a different WiSeNet might depend on the battery level, which is a feature of a single sensor node, not of the entire WiSeNet. So, it is difficult to define the behavior of a WiSeNet, which in general is the results of the behaviors of its sensor nodes.

### 4.2 Interactions between WiSeNets

Several challenges arise in the process of making different WiSeNets communicate.

Low-level communication, interference and power consumption. The first, obvious, constraint to build a NOWiSeNets is that the WiSeNet components are indeed networks, i.e. their nodes can exchange and propagate data. The second constraint to build a NOWiSeNets is that the WiSeNets share a low-level communication protocol. In principle, this protocol can be different from that used to exchange information between nodes belonging to the same WiSeNet. A possible starting point for both intra- and inter-WiSeNets communication is the IEEE 802.15.4 standard [9], which provides a common communication platform for devices that uses low power radios for low-bandwidth communications, as it is typically the case in WiSeNets. One of the challenges is to evaluate the suitability of 802.15 .4 to act as the low-level communication mechanism in the NOWiSeNets framework, possibly defining modifications of the 802.15 .4 communication protocols that better adapt to the peculiar features of the NOWiSeNets model.

An important aspect to be considered is how to deal with interference between spatially overlapped WiSeNets. This problem is likely to become more and more relevant as the number of deployed WiSeNets will increase in a near future. To deal with this issue, we can consider both single-channel and multiple-channel
solutions, identifying the solution that best fits into the NOWiSeNets framework.

Power consumption is another aspect to be taken into account when designing the node-level interaction mechanisms, as the efficient use of the limited energy available at the nodes is a major concern in WiSeNet protocol design. Power consumption reduction techniques can be applied at several levels, from physical node to node communication to high level applications.

Discovery and NOWiSeNets topology formation. Another requirement to build a NOWiSeNets is that the overlapping WiSeNets are capable to discover each other. The discovery protocol could be a mere extension of the one used by single WiSeNets to self-organize their topology at the time of deployment (in case of unplanned network deployment), or it could be especially suited to detect nodes belonging to other WiSeNets. Due to the interference problem mentioned above, this task might result very difficult and, in case a multi-channel solution will be adopted, it might imply the design of techniques for channel assignment, channel switching, channel scanning, and so on. A starting point for this challenge could be the current literature on neighbor discovery protocols [11], which, however, considers only the relatively simple case of discovering nodes belonging to the same WiSeNet.

Once two WiSeNets have discovered each other and established a (possibly rough) way to communicate, a number of steps have to be performed.

The node or the nodes involved in this inter-WiSeNet communication will be designed (or elected, in the case there are more candidates than necessary) as INs between the two WiSeNets. Since acting as interconnecting node could involve an increase in the node's burden, each WiSeNets could decide which nodes can act as INs, depending on the individual nodes' capabilities or on other dynamic factors, like residual energy.

One of challenges in designing node communication and discovery protocols is to allow interaction of mobile WiSeNets with other mobile or stationary WiSeNets. While in the case of two stationary WiSeNets we can easily imagine a sort of integration of the two networks after the discovery phase and INs election, in case of mobility the INs positioning or even WiSeNet's overlapping cannot be taken for granted in the long period. Thus, when designing the node-level communication primitives we could also consider ideas developed in the field of delay tolerant networking [12], which is relevant to our framework.

NOWiSeNets-level communication. Based on node-level interaction, we can devise communication mechanisms at a higher level of abstraction, where WiSeNets can be seen as autonomous, interacting entities. Communication primitives such as WiSeNet-
to-WiSeNet communication, broadcast, multicast, geobroadcast, and other primitives deemed relevant to the NOWiSeNets model could be designed. An example of geo-broadcast at the WiSeNet level is that of a smart home WiSeNet that sends an alert message to all the smart home WiSeNets in the neighborhood.

Another challenge is to define the mechanism used to route messages between far away WiSeNets. A number of decisions taken at higher level could have an impact at network level. For example, depending on mutual policies, two WiSeNets may exchange attributes declaring their will to use each other to relying packets. Or one WiSeNet could grant access to Internet to another WiSeNet. These decisions could clearly affect routing strategies in the case one WiSeNet has more that one way to perform his task. Indeed, in an ideal setting, we would like a WiSeNet be able to choose among different routing strategies (say, distance routing or geographical routing) in order to deliver data to, for example, an Internet connected gathering point.

### 4.3 Definition of NOWiSeNets application services

The definition of application scenarios makes it possible to identify some general types of services that interconnected WiSeNets can deliver to much wider groups of user. We indicatively mention five such general services.

- Connection service. This is probably the most obvious. A WiSeNet that is already connected to a fixed communication infrastructure, can act as a bridge for a disconnected one. The motivating scenario about chemical leaks fits also in this category. In fact, as already noted, the presence of a pre-existing network in the same area might represent a viable alternative (to the communication with the unmanned airplanes) in order to send real time data to the authorities.
- Alert service. This will be used in case of disaster events to promptly inform the WiSeNets in the vicinity of the potential danger. Another application of the Alert service is to inform the house owner who is traveling of a failure at home.
- Report service. This will be used to periodically get a summary of the observed phenomenon from one or more WiSeNets. For instance, a user might be interested in getting a daily report from the metropolitan WiSeNet that monitors air quality, to decide whether jogging is safe or not. Another example is that of a smart home WiSeNet, which periodically sends a report of the home conditions to the owner when she is traveling or working.
- Information Retrieval service. This will be used to retrieve any type of information (in principle), from the NOWiSeNet. For instance, a user might be interested in getting the measurements of air quality in the last
week from all the air quality check WiSeNets in a certain region. Note that, differently from the Information Retrieval services typically used in the Internet, in NOWiSeNet we must add a temporal dimension to user queries, as the measurements gathered from the physical world through sensors are time stamped. Implementing this type of temporal IR service implies the need of storing, at least for a reasonable time, summaries of the measurements taken by the WiSeNets. There is also the problem of defining a query language which is sufficiently powerful to express high-level questions [11]. Ideally, we would like to ask queries such that: "Which is the temperature of the bedrooms?", or "Does the lawn need more water?"
- Store and forward service. This service will provide facilities for storing records of specific events. As an example, a WiSeNet detecting the signal sent by a small radio device fastened to a seagull can store a record of the seagull transit and (later) send such record where it is needed. Scenarios like this could help improving the quality of monitoring animal behavior/conditions.

Besides the above, we envision a state of affairs where more application services will be required out of NOWiSeNets. All of this leads us to state the following research challenge:

Define new application services. Define new possible (and realistic) application services and implement them, as well as the five already listed, in the NOWiSeNets architecture.

### 4.4 Security and Privacy of NOWiSeNets

The realization of the NOWiSeNet framework will pose many security- and privacy-related research challenges. Two basic security mechanisms needed for interacting WiSeNets are authentication and access control. These topics have not been addressed previously for individual WiSeNets because either they are private networks or they are connected only to gateway machines with substantial resources that can implement standard userlevel access control and authentication techniques. Unfortunately, such standard techniques will be extremely difficult to apply in the NOWiSeNet scenario where access control and authentication must be implemented directly on the resource-constrained sensor nodes.

Developing a Framework for Access Control in NOWiSeNets. Clearly, users will require some basic access control before they allow their individual WiSeNets to interact with other WiSeNets. A relatively easy first step, which requires only the concept of unique WiSeNet IDs discussed earlier, is to have identifier-based access control. This would allow owners of WiSeNets to open access to specific known WiSeNets
with which they would like to interact. Of course, this access control mechanism will work as long as an unauthorized WiSeNet does not steal the identity of an authorized WiSeNet in order to gain access. Thus, this simple access control mechanism will also require authentication mechanisms for WiSeNets, which are addressed in the next paragraph. Access control mechanisms for WiSeNets should also permit different levels of access, e.g. access in order to relay packets through the WiSeNet, access in order to store data for later retrieval by other entities, and access to the actual data being generated by the WiSeNet. For each authorized WiSeNet, the exact level of permissions granted must also be set by the authorizing WiSeNet.
Inter-WiSeNet access control is complicated by the fact that each node in a WiSeNet is a possible IN. Thus, either the access control mechanism must be implemented within each sensor node or, upon each access request, an interconnection node must contact other nodes in the WiSeNet that implement a type of access control service. Issues of consistency of access control information and balancing of workload among the resource-constrained sensor nodes are two of many significant issues that impact the development of these access control mechanisms and which are unique to NOWiSeNet access control. Designing the structure and operation of an efficient and practical access control mechanism for NOWiSeNets is therefore a significant research challenge that must be addressed.

Developing Appropriate Cryptographic Mechanisms for NOWiSeNets. So far in this subsection, we have assumed that a WiSeNet knows the unique ID of another WiSeNet that wishes to interact with it. However, if there is some possibility of one WiSeNet masquerading as another, then strong authentication mechanisms are required. Simple authentication schemes such as passwords can be made stronger by encrypting authentication data using a secret key known only to the sender and the authenticator. This shared secret key can be pre-distributed in a secure way to the two interacting WiSeNets, or it can be generated dynamically and exchanged using public key cryptography. Stronger authentication schemes such as certificates digitally signed by trusted authorities require public key encryption techniques. Encryption is also needed for exchange of sensitive data, both between a user and a WiSeNet (possibly including relay of the data through other WiSeNets) and between two WiSeNets. Thus, what encryption capabilities can feasibly be provided on resource-constrained sensor nodes is an important research question, which will have an impact on the overall NOWiSeNets design.

Current research in WiSeNets assumes symmetric (shared) key cryptography, which has low enough computational cost to be implemented on sensor nodes [5].

It is assumed that some keys are pre-allocated to sensor nodes and the nodes use these keys for communicating sensitive data with each other when there is the chance of being overheard by an external entity [4]. For communication with the external world, it is assumed that the WiSeNets pre-allocated keys are also known to authorized users who do end-to-end encryption with the WiSeNet to protect data as it passes through intermediaries. Unfortunately, this assumption does not fit well within the NOWiSeNets paradigm. For example, in the case of a mobile WiSeNet, there is no a priori way to anticipate all possible WiSeNets that it can encounter and pre-distribute shared keys with all of them. Thus, in general, to enable authentication and secure communication between WiSeNets will require development of either lower-cost public key encryption mechanisms or entirely new ways of doing authentication and encryption specifically designed for the NOWiSeNets paradigm.

## Preserving Privacy while Maintaining Usability.

 In some cases, it might be useful for a WiSeNet to release some of its data to another WiSeNet that is not fully trusted but can provide a desired service. In these situations, the WiSeNet might wish to degrade the quality of the data being released (e.g. for privacy reasons) while still keeping the data in a usable form. Data obfuscation mechanisms [3], which are used e.g. in the area of privacy-preserving data mining [6, 7], could play a significant role in this regard. Obfuscation mechanisms are most effective when they are tailored to the specific type of data that is to be protected. Thus, development of effective obfuscation mechanisms for types of data generated by various kinds of WiSeNets is an important research challenge, the solution of which would facilitate privacy-preserving data exchange between partially-trusted WiSeNets.In conclusion to this section, it is our belief that developing primitives for security and privacy along the lines of those discussed above will enable higher-level operations to be carried out in a safe and secure manner. For example, in relating back to the motivating scenario described earlier, a WiSeNet must accept Alert messages (which cause, for instance, shutting windows, doors, and so on) only if the WiSeNet that issued the Alert has been properly authenticated and has the necessary permissions. Similarly, access to data about how many individuals are inside a house as well as their precise locations would be very important data to provide to a fire department in the case of a house fire, but it is critical that such data not be inadvertently released to a burglar who is surveying the property. It is our belief that the necessary security and privacy mechanisms can be developed in order to allow the full potential of NOWiSeNets to be realized.

## 5 Conclusions

In this paper we proposed the NOWiSeNets framework to fruitfully exploiting the inevitable interactions between WiSeNets that will occur in the next future.
It is our belief that, given its feature of enabling dynamic interconnection of WiSeNets, the NOWiSeNets paradigm has the potential to provide a further breakthrough to our ability of monitoring natural and human activities, and it can really be considered as tomorrow's Internet of the Physical World.

## References

[1] L. Anderegg, S. Eidenbenz, "Ad hoc-VCG: A Truthful and Cost-Efficient Routing Protocol for Mobile Ad hoc Networks with Selfish Agents", Proc. ACM Mobicom, pp. 245-259, 2003.
[2] L. Buttyan and J. Hubaux, "Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks", ACM/Kluwer Mobile Networks and Applications, Vol. 8, No. 5, October 2003.
[3] D. Bakken, R. Parameswaran, D. Blough, A. Franz, and T. Palmer, "Data Obfuscation: Anonymity and Desensitization of Usable Data Sets," IEEE Security and Privacy, 2(6):34-41, Nov-Dec 2004.
[4] H. Chan, A. Perrig, and D. Song, "Random key predistribution schemes for sensor networks," IEEE Symposium on Security and Privacy, 2003.
[5] L. Eschenauer and V. Gilgor, "A key management scheme for distributed sensor networks," Proc. of ACM Conference on Computer and Communications Security, pp. 41-47, 2002.
[6] W. Klosgen, "Anonymization Techniques for Knowledge Discovery in Databases," Proc. of the First International Conference on Knowledge and Discovery in Data Mining, pp. 186-191, 1995.
[7] L. Sweeney, "k-Anonymity: A Model for Protecting Privacy," International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, 10(5):557-570, 2002.
[8] http://www.xbow.com
[9] http://www.zigbee.org
[10] Frost \& Sullivan, "Wireless Sensors and Integrated Wireless Sensor Networks", Frost $\S$ Sullivan report, 2003.
[11] F. Zhao, L. Guibas, Wireless Sensor Networks: an Information Theoretic Approach, Morgan Kaufmann, San Francisco, CA, 2004.
[12] http://www.dtnrg.org


[^0]:    *School of ECE, Georgia Tech, USA.
    ${ }^{\dagger}$ Univ. di Modena e Reggio Emilia and IIT-CNR, Italy.
    $\ddagger$ IIT-CNR, Italy.

[^1]:    ${ }^{1}$ Of course, this is not the case of the WiSeNet owner, who in general wants to address specific sensor nodes in the network. However, the owner obviously knows the sensor node-level addressing mechanism used in his WiSeNet.

