Sensor Networks – Critical Infrastructure for Society? Challenges for Resilience, Security and Interoperability

Alistair Munro

EADS DS (UK) Ltd., The Quadrant, Celtic Springs Business Park, Coedkernew, Newport, South Wales, UK, NP10 8FZ
alistair.munro@eads.com

Abstract. Sensor networks have evolved rapidly in recent years from purpose built technology specific deployments to the stage where they are now able to deliver information through commodity communications channels. In parallel, pressures on the environment and resources have made the range of such purposes much wider and sensor input is necessary for sustainability of society's functions and infrastructure. The sensor network is becoming a part of critical network infrastructure and, in the same ways as other supporting technologies, must satisfy requirements for resilience, security and interoperability. This paper takes stock of requirements and solutions to highlight the opportunities and potential threats that follow from these trends. It reports on current work in CENELEC on a framework for addressing these issues.

Keywords: Sensor networks, resilience, security, interoperability.

1 Introduction

Sensor networks are widely installed in the world in urban, suburban and rural locations on the ground and in various airborne platforms, including balloons, high-altitude platforms (HAPs), unmanned airborne vehicles (UAVs) and satellites.

At present few of them have a purpose that involves real-time interaction with human beings. The Internet of Things will change this and make sensors, and actuators, first class devices, fully visible with end-to-end connectivity. We will depend on their capabilities and the data they provide for healthcare, energy management, environmental monitoring, transportation, homeland security and many other aspects of life.

Our assumption accordingly is that they are inevitably becoming a part of critical infrastructure. The motivation therefore for this paper is to explore the nature of, and challenges to, the processes that measure and control this dependency. The resilience of these systems will become a key discriminator of their quality and performance in achieving a positive or negative view of our reliance on them.

It is thus necessary to make sure that the sensor network applications operating over critical infrastructure are protected effectively as well as being thoroughly tested and correctly planned and deployed. If this is achieved then we can expect:

- Resilience of society's key functions:
- Improved situational awareness in anticipating and reacting to imminent events;

P. Chatzimisios et al. (Eds.): MobiLight 2010, LNICST 45, pp. 243–253, 2010.

- Better understanding of strengths, weaknesses, new opportunities, threats;
- Much more information available, so decision support is improved and reactions are higher quality;
- Systems are more efficient and cost-effective.

If it is not achieved then we risk:

- Dependency on systems that are not fit for purpose;
- Reduced security: less critical for disconnected systems, but essential when interconnected;
- Many kinds of attack: intrusion, denial of service, interception, masquerading;
- Poor interoperability – devices do not work together;
- Service level agreements not clear – the communications support may be inadequate or, at the other extreme, over-specified;
- Loss of privacy and confidentiality.

We begin with an overview of sensor networks in order to articulate the key issues in Section 2. The following sections focus on two key topics: Section 3 looks at vulnerabilities of the sensor communications system, and Section 4 discusses interoperability issues. Section 5 presents conclusions and topics for future study.

2 Sensor Networks – A Big Topic

As Akyildiz et al. have observed [1], significant opportunities have been created for systems composed of small, untethered sensor nodes, communicating over short distances using wireless media, by advances in micro-electro-mechanical transducers and actuators, wireless communications, and processing power of embedded digital electronic information processing devices. These are collectively termed wireless sensor networks (WSN).

Interesting problems arising from these opportunities have motivated extensive research into the communications aspects of WSN, of which resilience and performance is an important part, especially in terms of routing, data fusion and security. However WSN are not completely representative of a collection of technologies that is already well established and very widely deployed. Furthermore, the acquisitions of data is just the starting point for feeding it into the collection of applications that will use it. We will come to depend on these systems and the services they provide as part of society's critical infrastructure. Some illustrations of such systems and their vulnerabilities are given below:

- Electronic commerce: buying and selling, the banks and the financial sector. These are some of the biggest users of secure networked communications and they suffer persistent and sophisticated attack. Future e-commerce, online and on the move, will involve interaction of the financial systems with products and users to exchange identities, keys, and information about price, credit, user-guides and so on. Products will be equipped with RFID capability to store this information; and mobile personal devices will sense this information, exchange keys with local point-of-sale terminals, which will themselves interact securely on an open medium with the back-office enterprise system, e.g using UMTS.;

- The police, immigration, homeland security and other security services, the emergency services. This sector increasingly uses information from sensors, e.g. streaming CCTV, and other distributed services with an autonomous distributed machine-to-machine (M2M) component to gather intelligence from the environment. They require priority access to communications services, often needing to displace ordinary users and applications;
- Transportation, including highways, inshore water, railways, and civil aviation. The networking of transportation services is well established in certain infrastructure areas: toll collection by radio tag, traffic monitoring, signalling. Communication with, and between, road vehicles, boats and trains, including trams, is taking shape. There are requirements that are similar to those for civil aviation, e.g. collision avoidance. Wireless communications standardisation is quite advanced at the lower layers, e.g. CALM for road vehicles, or GSM-R for railways. Networking is largely focussed on IPv6;
- Resources, such as electricity, gas, oil, water and heat. The industry and governments are struggling to find a solution that will allow resources to be managed proactively, first for meter interoperability, the home sensor/actuator appliance network, and second for the communications network. There are interest groups, such as SmartGrids, supported by the EU Commission, and ongoing work in CENELEC in the Smart Meter Coordination Group;
- Environment, including: quality of air and water; disaster anticipation, first-response and recovery (fire, flood, earthquake, storm, attack – US Department of Homeland Security lists about 20 categories as well as civilian events). For monitoring, sensor networks are already deployed widely: wired and wireless terrestrial for a range of physical quantities on the ground; low to high altitude platforms, e.g. balloons, UAVs, for the lower atmosphere and short-range EO/IR, LIDAR and multi-mode sensing of ground state; and satellites for remote sensing;
- Health, including the enterprise, (public and commercial), telecare, and e-health. Experts from the healthcare institutions and from national and local government stress the demographic deficit that is heading our way: there will not be enough able-bodied people to take care of an aging population. Security is a particularly sensitive area: information harvested from sensors that was formerly private between a patient and a doctor now circulates on networks of various kinds.

While each of these sensor network applications has its own problems with security and vulnerabilities to attack, in general the resilience challenges focus on a two priority topics:

- Vulnerabilities affecting security and resilience at different layers of the communications systems;
- Interoperability of devices and their functionality from many suppliers and connected by heterogeneous media;

3 Security and Resilience of Sensor Network Systems

To highlight important issues, we will follow a layered model in the style of [1], and merge in security and resilience considerations during the analysis.

3.1 Media and Pathways

None of the pathways in common use are inherently secure or resilient. Even if they are not under attack, they can be disrupted in many ways, especially if the devices are in motion. Any kind of emission can be a clue that can be used by an adversary for intercept, jamming or masquerading, e.g. by replaying. If the attack is directed at a key pathway where traffic converges then the damage could be severe unless there is redundant capability: while more pathways give more opportunity for attack, the potential for data to travel by alternative routes may increase resilience. No deployment is static and, while there are specific issues for sensor systems, especially WSN, vulnerabilities of core and access networks must also be considered, as they too evolve: pathways that were placed in carefully selected locations and replicated to give resilience may move and converge, becoming key failure points.

3.2 Physical Layer (PHY)

While it might be assumed that the small sensors typically used in WSN are more constrained to narrowband, low bit-rate, short-range operation, the increasing maturity of systems that offer very low power, very high bit rate capabilities over even shorter distances, (UWB, or emerging 60GHz band wireless Ethernet), means that many new tradeoffs can be made taking account of a multiplicity of pathways that are difficult to intercept. At other extremes we can envisage a satellite with imaging sensors with very long paths visible to many potential intruders, or an appliance plugged into mains power lines that effectively broadcasts its messages to anybody connected to the power network.

The risk of instrusion and intercept increases according to how easy it is to demodulate and decode the energy into a bit-stream, so protective measures such as changing frequency, encryption keys or waveform can increase resilience provided the algorithms are not discoverable. Some of the PHY protective measures, e.g. evasion of attack by changing slots in a multiplexing scheme, may be done better by DLC/MAC functions;

3.3 Data Link Layer and Medium Access Control (DLC/MAC)

The systems that we are interested in have a networking requirement and the DLC/MAC layer provides essential support for this – it may itself have a range of routing and relaying functions that are in effect a network layer and provide for internetworking between clusters of nodes at DLC/MAC level. Being able to use one or more links allows techniques such as network coding, cooperative relaying and directed diffusion to be deployed. A data-centric sensor network may be operated entirely within this layer.

It is usually possible to identify a location in the communications system where sensors are logically, or physically, clustered. A proxy can be located at this point to interwork between the different DLC/MAC technologies. This could be a home DSL gateway or a set-top box; the GPRS systems used for meter reading applications are architecturally very similar but on a much large scale. These proxies are a point of attack from outside and from device inside the proxied network.

Complex systems are likely to have significant vulnerabilities, allowing an attacker to subvert their operation by attacking one or more DLC/MAC segments while apparently obeying all rules of normal operation. For example, an attack may be made by a device that is marginally hyperactive and ties up system resources by repeated re-registrations, or excessive traffic at critical times. Forwarding capability, such as routing or link virtualisation can multiply the vulnerability.

The risks of intrusion are substantial because there are so many ways of collecting packets. Sessions can be recorded, edited and replayed into the network without apparent intrusion or challenge. Commercial cellular communications systems are systematically secured to prevent access except to authorized users and, once admitted, their traffic is secure at least from each other on the wireless medium.

3.4 Network Layer

If there is a requirement for end-to-end connectivity between devices then we need to be able to establish paths across namespaces under different administrations, i.e. we need a naming and addressing scheme and a routing protocol. There is a large number of such schemes, e.g. as described in [4], some of which are IP-like internetworking systems using IP routing protocols for managed and ad-hoc infrastructures and others that operate using different principles, including data-centric approaches that require no routing protocol.

The class of applications that we outlined in Section 2 has a strong end-to-end, bidirectional interaction requirement in general and it is expected that they will use IP and its internetworking models in most, if not all, of such systems. Thus they will use routing protocols and there is potential for attack to subvert the routing relationships and the paths that data will take. Much of this can be done by sending false information. Popular examples of such attacks are:

- Sinkhole attack – the attacker node sends route packets with a low hop count value or other attractive metric value to its adjacent forwarding elements, e.g. the base station at the boundary between the WSN and the access network. The attacker will thus be able to alter the content of the data flow, throw it away, or launch additional attacks (e.g. selective forwarding attack, blackhole attack, and more);
- Replay attack – the attacker records routing traffic from genuine routing nodes and uses it to create a topology that may no longer exist;
- Wormhole attacks – similar to a replay attack but done in a different part of the network;
- Sleep-deprivation – the attacker generates spurious activity that will discharge its neighbor nodes and all nodes whose paths to the base station intersect the flooded path will have difficulty communicating with the base station;

Many managed systems will protect their routing traffic by securing the associations between routing nodes. However ad-hoc systems where every node is potentially a router are more vulnerable. Significant innovation is still needed to accommodate mobility at network level: the requirement may be that a device retains its network identity wherever it is attached (as is done in Mobile IP) but we must also consider mobile networks in cars and on people.

Proxy gateways are commonly used to mediate between IP and non-IP technologies. A proxy can be located at this point to emulate end-to-end IP connectivity or perform address mappings. This can be a place to attack.

As is well-known, the IP architecture is vulnerable to crude denial-of-service attacks on exposed network addresses.

3.5 Transport Layer

The transport layer is the extension of the communications service into devices and the processes they execute. Thus it understands end-to-end delivery and relationship with its peer device(s) at remote end(s), as well as the interaction with the processes at the ends. The quality of the outcome is a tradeoff of requirements for integrity and throughput, which is an application requirement, against the resources needed to achieve them.

The trend is towards increasing resources in sensor nodes so that it is possible to support the cost of a protocol such as TCP, or features built round UDP in the application to achieve reliable delivery. Attacks on transport-layer end points (ports) are familiar to IP users, so equipping sensor nodes with TCP or UDP will expose them to attacks commonly used on computers.

3.6 Application Layer (and the Others)

We include as aspects of the application the functionality of sessions, (e.g. TLS or IPSec associations), presentation (encoding of application semantics) and middleware for discovery and node configuration, e.g. UPnP, or key-exchange and other supporting functions. All these protocols expose new information about sensor nodes (those that are capable of using them) or the proxies that implement them on the sensors' behalf.

If an intruder is able to establish connectivity at such a deep level in the system and its nodes then protection mechanisms that are supposed to prevent this have failed. This is anyway a contingency that must be anticipated: no system is perfect and its users and administrators will make mistakes or act maliciously from time to time. Maybe the measures that were provided are very simple and present only limited barriers: maybe we want to attract attackers and encourage them to give themselves away.

Depending on the application, the attack may have impacts ranging from none to catastrophic. The intruder may replay past disruptions, or create situations that appear plausible but do not exist. To achieve resilience, the system should be able to audit traffic (its originator, destination and route) and the assets connected to it. These are difficult and expensive to do, and it is inevitable that there will be a certain level of noise and interference affecting any information.

3.7 Current Approaches

As well as recent ongoing research work, e.g.[2], [3], [5] and [6], relating to intrusion detection and security in sensor networks, there are several standards that have been written to categorise security vulnerabilities and threats and define the functional capabilities to counter attacks. For example, from the ITU-T there are X.800 and X.805 that cover the larger network for end-to-end communications, and from ISO/JTC1/SC6 there is ISO29180 (currently at CD status) giving a framework for

security in sensor networks. Home network security is described in ITU-T X.1111 – X.1114. Overall these standards reflect a model of the security issues split between the external networks and internal ones.

The specific vulnerabilities of a sensor network are described in ISO29180, including:

- Many nodes, but not all, lack the capability to implement security mechanisms for well-known functions. It may not be possible to use public key systems. The sensor network may be especially vulnerable to DoS attacks based on functions related to key generation and exchange;
- Compromise of nodes, when attackers evade security measures and gain access: possibly through tampering to connect directly to the electronics of the sensor; or by being able to connect to the sub-network or route via external networks to communicate with and subvert the function of the devices. Such compromises may happen because of faults in the application, interoperability failures, or poor system design: nodes that accidently make an attack may not be aware that they are behaving badly;
- Evolving deployments. The configuration of the initial deployment of the sensor network may not be known if it is scattered randomly, or it may be systematically installed and documented. A given configuration will change when sensors change position, network connectivity, or have their functionality enhanced (or reduced) or upgraded. Faults will also change the configuration, maybe transiently or permanently: loss of connectivity through lack of coverage or jamming; loss of power; and simple failure;
- Key failure points when traffic flows through a single device, such as a home gateway or a fusion node that has become the focus. This can be avoided, at a price, by exploiting redundant paths made accessible through the locally available media.

ISO29180 also identifies attackers from inside the sensor network and from networks that connect to it. The specific threats that these present include, (with reference to [5]):

- Destruction of, gaining control of, and tampering with, information or device capabilities, (hijacking);
- Intercept and disclosure of information, (eavesdropping);
- Generation of information that is incorrect (semantic disruption);
- Disruption to services, in particular to routing.

These reflect the vulnerabilities noted above, and the risks will be affected by the extent to which the sensor devices and supporting gateways are able to counter these direct threats. Disrupted routing is an especially serious threat, particularly so when an attacker can place itself at a forwarding point where it can change the pathways in addition to the threats mentioned above.

In the context of critical infrastructure, the risk assessment of the threat of intrusion must be realistic and strict. It is not likely that a separate physical core network will be installed for our family of sensor network applications: sections of the access network may however be physically separate, e.g. sensors at trackside or roadside follow the railway and highway maps. The technology for segregating traffic in a shared

network using virtualization from data link layer upwards is well understood. However, an attacker could deliberately connect the infrastructures together. If routing protocol flows across this connection then the apparent topology of any of the involved infrastructures could change and traffic would mix and flow in unexpected ways. Self-organising applications, that discover and configure devices and functionality without user intervention might enroll inappropriate functions or other attacking nodes, and thereby cause unwanted behaviour. A person's home in which telecare and energy management applications coexist is exactly such an interconnection point, and one that would be expected to exist.

4 Coexistence, Interworking and Interoperability

Because multiple applications, services and devices will share resources and infrastructure to some extent, it is essential that they do so without conflict. Additionally, it is certain that communications systems will become more heterogeneous, not less, especially where wireless technologies are used. Thirdly, there is already a wide diversity of standards in sensor network domains. These standards come from a range of bodies – the IEEE, the ITU, ISO, ETSI, and the IETF – and may fulfil the same function in different ways. Some are regional or sector specific even if they have the status of International Standards. This diversity is especially apparent where bodies traditionally concerned with professional ICT equipment (the IEEE) or telecommunications (the ITU, ETSI) are developing specifications for sensor systems or machine-to-machine communication to be used by consumers in general.

The presence of multiple processes active in the same space, physical or electronic, leads to failures of interoperability. This is a term that is widely used and is understood in different ways: it is more than just the ability to communicate and exchange bits across a collection of technology-specific pathways. The bits must be combined and formatted as structured data; the data must be comprehended as information independent of representation; and the information must be used and acted upon in a consistent way to achieve real effects. However, the term is used in all these contexts and many others, e.g. the rules of working between the police, fire and ambulance services, so our interpretation must be clearly stated.

In managing the solutions to interoperability problems, we make a distinction between aspects that are mainly technical and those that are concerned with processes outside ICT and electronic domains of sensing and actuation, e.g. the management of a domestic heating system to an occupant's requirements, taking account of the weather and the heating cycle of the building. Here we are interested primarily in the former. The technical aspects of interoperability should be uniform and consistent for all processes so that the ambition of sharing infrastructure and resources can be achieved.

At a technical level, we can express three requirements, each of which has an interoperability aspect:

- Co-existence - where different systems can operate in the same environment without hindering each others' operation or otherwise conflicting: e.g. a home wireless security system using Zigbee, a WiFi local area network, and a Bluetooth telecare service. All occupy the same ISM spectrum and will potentially interfere with each other but they are separate applications involving interoperable devices;

- Interworking - where different technologies are connected together to transfer data end-to-end. It is primarily a technical solution encompassing connectors, protocols, bridges, etc. An example is where the home security system above is connected to a PC monitoring application using the WiFi network to transfer the data between the security system and an external web application. At some point the different communications systems come together at one or more gateways and information can flow between them;
- Interoperability is where different application functions interoperate with each other: this adds business rules, processes, security provisions, etc. that enable applications to be joined together and to use devices of any provenance. The example here is the home wireless security system being controlled and monitored remotely using a separate web application, possibly sending alerts to the owner's mobile telephone or, locally, to the TV set.

The new challenge for interoperability is to ensure that the three requirements are met for systems of variable structure and population; changed and evolved by their users, who will generally be non-professional consumers lacking technical expertise; operating distributed algorithms and protocols with a significant level of machine-to-machine autonomy; as well as executing a range of applications, which may change from time to time, sharing resource, devices and infrastructure.

We have seen already, in section 3, how this variability impacts on security and resilience. Let us now look at it from a specific industry and standardization viewpoint.

4.1 Current Approaches

One area where sensor networks are likely to be widely deployed is the home, e.g. for management of resource consumption (electricity, gas and water), telecare, and a wide range of personal applications, such as entertainment, security, or comfort control. These are termed collectively Home and Building Electronic Systems: HBES.

There are already many standards established for the communications systems supporting HBES: IEEE 11073 for healthcare devices, EN50090 (CENELEC TC205) and EN14908 (CENELEC TC247) for control and communication. There are several activities in ETSI (the M2M group), ITU-T (e.g. G.hn and G.cx, or the G.6690 recommendation), and ISO/SC25/WG1 that have a direct impact on future directions. Where HBES interact with other systems, a range of ICT standards (e.g. UPnP) or IP protocols are used.

Because of the diversity of solution and the large number of legacy deployments, agreement on a single specification or standard is unlikely, so the need for an interoperability solution meeting at least the three requirements outlined above is a priority. Even this is contentious: maintaining interoperability even for devices implementing the same standard has proved difficult, so what prospect of success can be expected when multiple systems are interconnected?

To begin to address this problem in the light of the challenge posed by highly dynamic systems. CENELEC TC205 approved in July 2008 the formation of a Workshop that will deliver an Agreement that defines an Interoperability Framework and requirements for complying with it. Such an Agreement is voluntary and will be developed into a collection of mandatory conformance requirements in the form of a CENELEC Technical Specification. The Workshop is close to finalizing the text of the Agreements, which has the following key points:

- A categorization of Interoperability Levels, as shown in Table 1. This associates coexistence, interworking and interoperability between dissimilar systems with levels 1, 2 and 3, reflecting the current state of the art. Levels 4, 5 and 6 look into the future to highlight the greater dynamics of HBES systems, the prevalence of machine-to-machine autonomy at local and remote level, and the wider range of applications that will reach into the home;
- A categorization of middleware functions needed to support the formation of applications at level 4 and above: device discovery, application configuration, and management. A survey of the 17 or so contemporary standards in the HBES domain showed that these three groups of functions were provided by all of them in isolation, even though the way they work varies in detail;
- A set of conformance requirements placed upon devices, (sensors, actuators, gateways/routers, and the controlling application components) to enable them to state compliance with the framework of levels and functions, and identify the objects that they implement;
- Specific provision for security and measures to enhance resilience following the models of section 3 above;
- A formal specification of the structure and required content of an Interoperability Implementation Conformance Statement to be applied to products regardless of their provenance and the standard(s) that they implement.

Table 1. Interoperability Levels

0	A single system of supplier-defined structure built from devices using a single HBES specification and locally defined interoperability verified by the supplier for one or more application domains. No assurance of coexistence is provided.
1	A Level 0 system.operating across one or more application domains Verified coexistence is required.
2	Multiple Level 1 systems that interwork to exchange information and interoperate across specification and application domains verified by the suppliers using conformance specifications.agreed by each HBES specification used.
3	As Level 2, and the interoperability is verified with respect to international standards applicable to the HBES specifications used in the system.
4	As Level 3, but conforming to IFRS so that the applications and devices can be installed, managed and changed during the operation of the system by a qualified installer.
5	As Level 4, and changes of application and devices will be done automatically.
6	As level 5, and with remote management, diagnostics and maintenance. (automatic installation, operation and support).

An assurance of interoperability end-to-end between devices and application elements is an important complement to measures taken to ensure quality and grade of service of the network and security of its operation. While the CENELEC Workshop Agreement is a starting point in achieving this assurance, it is likely to take some time before it becomes acceptable as a Technical Specification and mandatory.

5 Conclusions and Future Work

The work reported in this paper has considered the security, resilience and interoperability issues that will be important in fulfilling expectations of quality of service, availability and integrity of sensor network applications that will form part of the critical infrastructure of future society.

This infrastructure, and the resources and devices connected to it, will be shared by many applications executing on behalf of millions of people, at home and on the move, and involve a multiplicity of services in the home, at work and in public spaces. Several issues remain to be addressed, related to the scale in population and diversity of behaviour:

- Integration of very large numbers of sensors – how much, or how little, state must be stored to maintain connectivity, implement effective sharing, record usage, and ensure security;
- Unusual traffic distributions triggered by human activity or the environment, or autonomously between machines, with interaction occurring at machine timescales. These could be an indication of attack, or maybe just unexpected behaviour;
- Managing connectivity and presence of a large number of small networks, fixed and mobile;
- Developing a formal model of interoperability as support for greater resilience and security.

Acknowledgement

The author is grateful to ENISA and TAHI (The Application Home Initiative) for partial support in doing the work reported here.

References

[1] Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E.: Wireless sensor networks: a survey. Computer Networks 38, 393–422 (2002)
[2] Braun, T., Danzeisen, M.: Secure Mobile IP Communication, Uni. Bern internal report
[3] Anand, M., Cronin, E., et al.: Sensor Network Security: More Interesting Than You Think. In: Proceedings of HotSec 2006, Usenix Workshop on Hot Topics in Security (2006)
[4] Karaki, J.N., Kamal, A.E.: Routing Techniques in Wireless Sensor Networks: A Survey. IEEE Wireless Communications, 6–28 (December 2004)
[5] Martynov, D., Roman, J., Vaidya, S., Huirong, F.: Design and implementation of an intrusion detection system for wireless sensor networks. In: IEEE International Conference on Electro/Information Technology, pp. 507–512 (2007)
[6] Mitrokotsa, A., Karygiannis, A.: Intrusion Detection Techniques in Sensor Networks. In: Wireless Sensor Network Security. Cryptology and Information Security Series, vol. 1, pp. 251–272. IOS Press, Amsterdam (April 2008), ISBN 978-1-58603-813-7