

Foresight

Infectious Diseases: preparing for the future

OFFICE OF SCIENCE AND INNOVATION

**S3: State-of-Science Review –
Intelligent Sensor Networks**

Professor David De Roure
Dr Danus Michaelides
Dr Mark Weal

School of Electronics and Computer Science
University of Southampton
dder@ecs.soton.ac.uk, 023 8029 2418

This review has been commissioned as part of the UK Government's Foresight project, Infectious Diseases: preparing for the future. The views expressed do not represent the policy of any Government or organisation.

Contents

1 Introduction.....	3
2 Background: The evolution of computing	4
2.1 Hardware perspective	4
2.2 Software perspective	4
2.3 Information perspective.....	5
2.4 Summary	6
3 Definitions: What is an ISN?.....	6
3.1 Sensor	6
3.2 Sensor nodes.....	7
3.3 Sensor networks	7
3.4 ISNs.....	8
4 Characteristics of ISNs	9
4.1 Power budget.....	9
4.2 Durability.....	9
4.3 Resilience	10
4.4 Reliability	10
4.5 Mobility.....	11
4.6 Autonomy.....	11
4.7 Scalability.....	12
4.8 Heterogeneity	12
4.9 Security.....	12
5 Exemplar ISNs	13
5.1 FloodNet	13
5.2 Pollution monitoring	14
5.3 Well-being monitoring	15
6 Lifecycle of ISNs.....	15
6.1 Device design	15
6.2 Deployment.....	16
6.3 Configuration and maintenance	16
6.4 Data processing and transmission	16
6.5 Information processing.....	17
6.6 Notification and information	17
6.7 Information re-use.....	17
7 Summary	17
7.1 Key summary points	18
7.2 Future trends	19
Acknowledgements	20
References	21

1 Introduction

Intelligent sensor networks (ISN) perform the collection of reliable and timely information which in turn enables early warning, supports decision making and helps provide rapid and co-ordinated responses to potential threats. This chapter is a state-of-science review covering aspects of intelligent sensor networks which are relevant in the context of the detection and identification of infectious diseases (DIID).

Sensor networks are a powerful tool in DIID. They are developing rapidly, enabled by the increasing miniaturisation and decreasing cost of hardware, together with innovations in sensor and multi-sensor techniques, networking technologies and data processing. We anticipate a deployed ISN infrastructure for DIID which assists in screening and can be coupled with a rapidly deployable infrastructure to gather more data or assist in responding to situations.

We are moving into a world where increasing numbers of computing devices are deployed on our person and in our environment – the world of pervasive or ubiquitous computing (Weiser 1991). These products and deployments, augmented with specialised facilities as appropriate, will provide a basis for a DIID infrastructure. We already observe the extensive adoption of consumer electronics products, such as mobile telephones and portable digital assistants (PDAs), which effectively constitute a pervasive computing infrastructure. Future devices will be sufficiently small that they can be deployed in novel ways, as illustrated by the ability to implant devices in the body or to disperse devices in a physical environment.

Relevant practice in sensor networks is becoming established, for example, in monitoring well-being (De Roure 2005; Laursen 2006), and intelligent techniques are also emerging as exemplified by the prioritised monitoring of environmental conditions (De Roure 2005).

In this review, we focus on the capabilities of ISNs. The techniques used by the actual sensors in the DIID context, and the corresponding data fusion and data mining techniques, are discussed elsewhere.

The review is structured as follows. In the next section we step back and look at the bigger picture, taking account of other trends in computing such as grid computing, agents and the semantic web – this puts the evolution of ISN in context. We then describe what is meant by an ISN in Section 3. The technical characteristics of an ISN are presented in more detail in Section 4. Recognising the importance of focusing on systems, we look at exemplars in Section 5 and consider the lifecycle of an ISN in Section 6. We close in Section 7 with a summary of the main points arising from this review.

2 Background: The evolution of computing

2.1 Hardware perspective

It is clear that computing devices are becoming smaller and cheaper. Underlying this is the famous prediction of Moore's Law, that the number of transistors on a chip will double every 18 months. Moore's Law also explains the continuing increase in computing power per unit cost.

There are other significant hardware trends. Computers also contain hard disk storage and RAM memory. The capacity of a hard drive has been increasing more rapidly than transistor counts. RAM storage capacity is also increasing, though it is limited by speed in comparison to processor speeds. Many new data transmission methods are possible. There are wireless communication techniques, including the now familiar Bluetooth and cellular solutions, and also passive techniques such as radio frequency identification (RFID) tags.

The implications of these technology trends are more than just smaller or faster computers. We are seeing a very significant change in how these new technologies are deployed and used. For example, the technologies are being diffused into everyday objects and contexts, with the promise of supporting and enhancing people's lives. With this comes a fundamental shift in how we design, use and think about computers, and in the structure and flows of information.

Behind the scenes there is now considerable computational and data processing power, whether through supercomputers, piles of commodity PCs or spare capacity on PC installations. Today's networks have higher bandwidths on the wide area than those used locally, in complete contrast to a few years ago. Optical networking, for example, is resulting in a paradigm shift in how we tackle data processing in the wide area (Smarr et al. 2003).

In fact, the increasing number of devices that collect data demand increasing processing capabilities for that data. For example, the increasing volume and density of data gathered by sensor networks demands greater processing, especially with respect to data analysis, modelling and prediction. The increasing ability to work with the data globally is relevant to the DIID context. It is clear that on the timescale of this study – 10–25 years – there will be significant changes in our notions of computers and computing.

2.2 Software perspective

Hardware aside, software techniques are also adapting to this changing deployed infrastructure and context of use. The challenges result from working with increasingly large-scale, heterogeneous systems that are dynamic and flexible. This is resulting in innovation in key respects:

- dynamic virtual organisations. Resources – both computational and people – will come together in flexible arrangements to meet the requirements of the task at hand. This is exemplified by emergency response scenarios.

- self-configuring, self-managing, self-optimising, self-healing systems. The components of these large-scale systems will not be configured by hand, but rather the system will have so-called 'self-' or autonomic capabilities.
- information and knowledge processing. The pervasive intersection of the digital and physical worlds is delivering vast amounts of data for analysis, mining and re-use, in both anticipated and unanticipated ways.

Several technologies are being brought to bear on these challenges, including, for example, grid computing (Foster et al. 2001) and a combination of grid and agent-based computing (Foster et al. 2004). All these challenges require some notion of 'machine-processable descriptions' of objects, services and resources in order to support automation. The combination of grid and semantic web (Berners-Lee et al. 2001) technologies is known as semantic grid (De Roure et al. 2005).

Predictions about the evolution of grids suggest the notion of the service-oriented knowledge utility (NGG3 2006). This approach to service-oriented architectures and delivering knowledge may well provide an appropriate architecture for aspects of DIID.

2.3 Information perspective

It is not sufficient to consider the sensors or devices in isolation. An information perspective is essential, as fundamentally this is what DIID is about.

For example, measurements are not the only information associated with a sensor. Contextual information (such as location, age, calibration parameters, etc.) about a sensor is crucial to the interpretation of signals from that sensor. Furthermore, provenance information is required to interpret data, understand its quality and accuracy, and to determine origin and enable trust (Wong et al. 2005).

Tracking diseases (loss of wellness) in animals with regular habits such as territory marking or confined ranges (Laurson 2006) is relatively easy as one can often do the trend analysis based on a single collection point (or a small localised set). Data fusion is needed when the data is collected in diverse locations, for example, at migration sites. This is a good example of the need for a datagrid in order to store and process the collected data.

There are also interesting social trends in how information is generated. Traditionally, there was a clear distinction between content providers and their consumers. Now we see consumers empowered as content or service providers and a model of content creation that is based on participation and sharing. For example, Intel's Place Lab system (Hightower et al. 2005) enables laptops, phones and PDAs to estimate their location based on radio beacons such as wireless access points, and the information about the locations on the beacons is also mapped by the community. This evolving ecosystem can prove very powerful. Self-collection of data raises particular issues of trust and provenance.

2.4 Summary

ISNs reflect all the trends discussed above. In 10–25 years we will be working in the context of a vast, heterogeneous pervasive computing infrastructure which itself will be a self-organising ecosystem. The sensors represent the intersection between the physical and digital worlds – they are tangible artefacts. The grid and semantic processing that occurs behind the scenes is invisible but also crucially important in order to be responsive in the face of the data deluge created by sensor networks.

We discuss these technologies to the extent that they support the ‘intelligence’ of the sensor networks. However, the task of data fusion from multiple sensors, and the task of taking a large volume of data and extracting some meaning from it – often called knowledge discovery – are also crucial parts of this picture and are addressed in the separate review on Data Mining and Fusion.

3 Definitions: What is an ISN?

3.1 Sensor

A sensor converts a real-world measurement – such as a biological response in the case of a biosensor – into an electrical signal.

In general, uses for sensors might include environmental and habitat monitoring, biosensors, condition-based equipment maintenance, disaster management and emergency response. Simple sensors might be used to detect information about the environment: temperature, humidity, lighting conditions, soil make-up, noise levels, carbon monoxide levels, the presence or absence of objects, chemicals, radiation, enzymes, etc. Sensors may also require specific resources to carry out their task, for example, specific chemicals to detect the presence of foreign agents. Tracking the location of objects, whether they are physical assets or animals, itself provides information that supports the process of DIID.

The position of the sensor might be a crucial important factor, and knowledge of its physical context may be essential to the interpretation of the output from the sensor. This is, for example, the case with sensors monitoring physiological signals on mobile patients, and in environmental monitoring through sensors that have been dispersed into the environment or are mobile due to their carrier, e.g. attached to animals. Multiple sensors may be used to gather this information, for example, through geographical positioning systems (GPS) to identify location and accelerometers to record movement. RFID tags provide a means of electronically identifying particular physical objects relevant to the interpretation of sensor readings.

We draw a distinction here between remote sensing and what we will term *in situ* sensing. Remote sensing can be defined as a technology used to acquire information about an object by detecting energy reflected or emitted by that object when the distance between the object and the sensor is large.

Typically, it is the gathering of data about the Earth and the environment from satellites. This is beyond the scope of this particular review. The integration with *in situ* sensing has been discussed elsewhere (Teillet et al. 2001).

We will focus on *in situ* sensing, the acquisition of information when the distance between the object and the sensor is much smaller, i.e. the sensors are deployed inside the phenomenon or very close to it. Where measurements or observations are made from nearby locations that are not strictly speaking *in situ*, the expression 'proximal sensing' has also been quite widely adopted.

3.2 Sensor nodes

In this review, we refer to a 'sensor node', which in the context of an ISN is typically a sensor (or sensors) coupled with hardware and software for information storage, processing, transmission/reception and to provide power. For example, today a mobile phone or PDA equipped with a sensor can be thought of as a sensor node, as is a purpose-built portable instrument or a laboratory instrument with these capabilities.

Sensor nodes may be part of hand-held equipment, embedded in our surroundings, in the natural environment or other artefacts. They may combine multiple functions. Consider, for example, current mobile phones which have audio, image and video recording, storage and replay capability, can communicate live or via a store-and-forward model over cellular networks, Bluetooth and infrared, and can store perhaps a gigabyte of data – they can be regarded as sophisticated sensor nodes.

Sensor nodes are a subset of the diverse range of pervasive computing devices that will surround us. Other devices may incorporate visual displays, audio or other mechanisms, or effect a mechanical operation. Some may have a role as intermediary in data processing and transmission without having a sensing or actuating capability themselves.

As processing power increases for a given size of device, more storage and computation is possible locally (as long as there is sufficient electrical power to do this). This increasing computational power will facilitate local analysis of data.

3.3 Sensor networks

A sensor network is an interconnected set of sensor nodes, gathering information in a co-ordinated way (Estrin et al. 1999). By having a network of sensors, sampling can take place over an area and the ability exists to exploit redundancy and dynamic communication mechanisms between the sensors. This is important because sensors cannot be assumed to be reliable or always available.

It is this communication that forms a key component of much sensor network research. Not all sensors in a network may be in a position to communicate their information directly to the information receiver (sink). Figure 1 illustrates

a basic sensor network. The sensors form a sensor field, roughly denoting their area of coverage. Individual sensor nodes pass and receive information through the network with some nodes communicating with the sink. This in turn can relay the information by a variety of means (Internet, satellite communication) to processes that wish to deal with the sensor data. Strategies adopted from networks and distributed systems research allow the communication of the information via other nodes in the network.

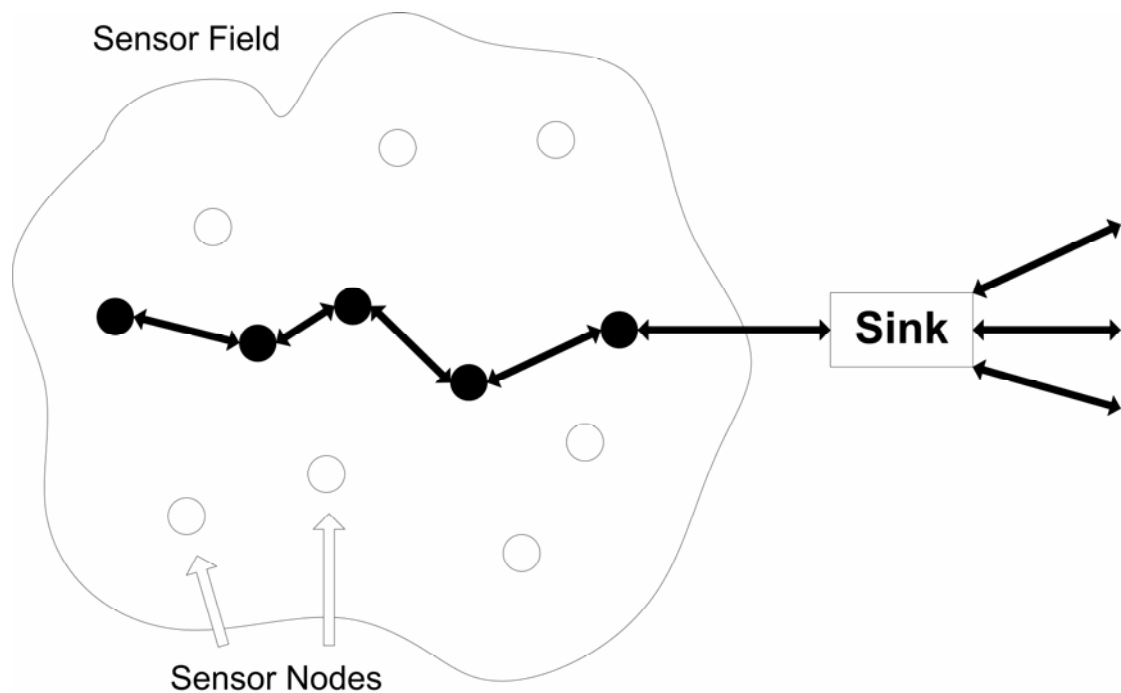


Figure 1: A sensor network

3.4 ISNs

What makes a sensor network intelligent? An ISN is one that can adapt to changing situations and present information that is relevant to the current circumstances.

The intelligent behaviour might include reasoning, adapting communication strategies, intelligent monitoring and expenditure of resources, coping with failure and topology changes as well as the ability to learn. An ISN may not have all of these behaviours, but a degree of autonomy would differentiate them from simpler sensor networks.

The network may also have some level of self-organising capability, both on initial deployment and as the network changes over time through failure, movement of sensors or environmental factors. Co-operation between nodes is another area where intelligent networks can achieve more than simple deployments. The ability to carry out simple computations locally and transmit

only the required and partially processed data is also desirable, along with the ability to decide when it is best to compute and when to communicate.

4 Characteristics of ISNs

Extracting data gathered by sensor nodes in remote locations involves some unique challenges. In this section, we will look at some characteristics that are common to many intelligent sensor networks and are key subjects of research and development.

4.1 Power budget

Energy efficiency is often a crucial problem in sensor networks. Sensor nodes will often carry limited, generally irreplaceable, power sources. Typically, sensors will have power requirements for the three general tasks of sensing, communication and data processing, although specific sensors may have additional requirements – perhaps actuators are attached or location identification is also required. Node-level techniques and network-level techniques are employed for various efficient energy management methods. Long-term operation of sensor networks requires economic power management (Shah and Rabaey 2002).

Traditional networks aim to achieve high quality of service provisions. With sensor network protocols, however, the focus will primarily be on power conservation. This will involve in-built trade-off mechanisms, potentially with the option of prolonging network lifetime at the cost of lower throughput or higher transmission delay. Dynamic power management schemes for sensor networks have been proposed that provide a variety of modes of operation (Sinha and Chandrakasan 2001).

Time scheduling is another technique used to minimise power consumption, reducing and co-ordinating duty cycles. Optimal behaviour of a sensor network may depend on current prioritisation of data capture which may in turn depend on the actual data – this kind of analysis and response is an aspect of intelligence.

Technology is now available to harness energy from mechanical vibration, which has considerable potential for powering sensors in certain situations (Glynne-Jones et al. 2004). Energy-aware sensors are also an important subject of research, especially as our deployments move from a smaller number of larger devices to large numbers of smaller devices. By way of example, the MIT μ AMPS (Adaptive Multi-Domain Power-aware Sensors) project is developing a framework for implementing adaptive energy-aware distributed micro-sensors (Min et al. 2002).

4.2 Durability

Sensor networks are often deployed with the expectation that they will operate for long periods of time. This could be because their placement prevents ongoing maintenance or reconfiguration or perhaps simply that the initial

deployment effort is large. It is often desirable therefore that the network has long-term durability. This applies to the individual nodes and the network as a whole.

The sensor network, and the sensors themselves, may need to withstand particular (possibly hostile) environmental conditions such as temperature, pressure, vibration, bio-fouling and chemical damage. To illustrate an extreme case, in the GlacsWeb system, the probes are deployed at the ice-sediment boundary 50–80 metres below the surface of the glacier (Martinez et al. 2004). Each probe is equipped with pressure, temperature and orientation (tilt in three dimensions) sensors and has to cope with the pressures and low temperatures within the glacier. The nature of the deployment also means that the probes are not recoverable, increasing the need for durability as replacement or maintenance is not possible

The network must also be durable in respect to its usage of resources. Aside from power, discussed above, sensor nodes may have other resources that are depleted over time. If the sensor is testing the presence of organisms in its proximity by using a particular chemical, it may have a limited quantity for its use, i.e. it can only make a limited number of samples. Deciding when to take samples may require strategic decisions across the network so as to maximise the durability of the system as a whole and promote the longevity of the network's sensing ability.

4.3 Resilience

Sensor nodes are prone to failure and the network must be resilient to this. This is especially true where nodes are being deployed in hostile environments. Individual sensors may fail through lack of power, have physical damage, or suffer from environmental interference. The network as a whole has to be resilient to these challenges. This form of resilience is often described as fault tolerance (Shen et al. 2001), the ability to sustain network functionalities despite sensor node failures.

4.4 Reliability

Reliability in intelligent sensor networks can apply to both the reliability of delivery of sensor information and also the reliability of the information itself. Sensor design will obviously have an impact on the reliability of the information produced. Self-diagnostic sensors, sensor redundancy and detailed capturing of provenance can all be used to improve the reliability of the information received. It is important to capture as much information as possible in order to interpret the data, including understanding the accuracy. A familiar example is that GPS data also carries information about the number of satellites, and it is useful to record this in order to properly interpret the positional information.

In some cases, it will be important that the sensed information is received reliably, i.e. the network behaves reliably. Research from the field of networks on quality of service can be applied here and use of *ad hoc* networking

techniques to enable sensors to dynamically reroute their information in order to keep the data stream alive.

4.5 Mobility

The position of the nodes need not be engineered or predetermined. This allows random deployment in inaccessible terrains, which, depending on the application, may provide useful opportunities, such as air-dropping sensors, attaching sensors to animals, or sensors being carried by individuals.

Mobility need not refer to initial deployment, however. During its lifetime of operation, the sensors themselves may move around. This may be a designed factor (the sensors are attached to animals that roam a wide area) or unintentional, due to earth movements or environmental factors. There are different strategies that can be adopted to allow the network to cope with the movement of its component parts. The overall network needs to be able to smoothly adapt to the changing topology of the network of sensors.

A final aspect of mobility concerns redeployment. With sensor networks that are intended to be redeployed, the mobility of the system becomes essential. It may be a requirement that a sensor network deployed to monitor the spread of a contagion can be quickly collected as one epidemic is controlled and redeployed in areas where the epidemic is suspected to be spreading. This redeployment should ideally be rapid and with minimum reconfiguration following the movement of the sensors.

4.6 Autonomy

The question of 'what to compute' and 'what to communicate' is critical to sensor network behaviour. Increasingly complex sensors are able to do more processing locally and apply intelligent processing within the network. What is then communicated out of the network can be a summary of the information. Trade-offs need to be made, however, once the data has been extracted from the network, processing can be relatively cheap, take advantage of technologies like the grid and, importantly, not use up power resources from the network. The communication of data out of the network does have power implications, though, so reducing the volume of information transmitted by applying on-sensor processing can be important.

DIID sensor networks may typically be vigilant in a screening role and then when certain events occur their behaviour will change, perhaps requiring better power.

From the perspective of those responsible for subsequent processing of the collected data (such as data mining), choosing to discard anything at the sensor is to be avoided as far as possible. It may be possible to communicate urgent data but store historical data.

4.7 Scalability

DIID systems will consist of a vast number of devices. In some cases, the accuracy of data will be the result of large numbers of partially redundant measurements, in other cases it will be the high-quality measurements from specific devices.

It is necessary to regularly add groups of sensors as well as manage potentially large numbers of systems. Unlike *ad hoc* networks, the number of nodes in a sensor network will often be several orders of magnitude higher than in an *ad hoc* network deployment. The nodes can also be very densely deployed whereas *ad hoc* networks are usually designed to maximise coverage with the minimum number of nodes. Hundreds of thousands of nodes might be deployed within tens of feet of each other and scalable communication paradigms are essential (Intanagonwiwat et al. 2000).

4.8 Heterogeneity

It is inevitable that sensor networks will become increasingly heterogeneous. Future deployments will need to incorporate legacy systems to maximise their utility. Networks will also comprise a variety of sensors, some providing usefully contextual information about other sensors which can be used to improve interpretations of the data produced.

The BLOSSOMS (Building Lightweight Optimised Sensor Systems on a Massive Scale) sensor network project aims to identify research issues at all levels from practical applications down to the design of sensor nodes (Gao et al. 2004). The project is developing a heterogeneous sensor array, including different types of application-dependent monitoring sensors and intruding sensors. Application-dependent power-aware communication protocols are also being studied for communications among the heterogeneous sensor nodes.

4.9 Security

All sensor network levels may be required to take security into account. In some cases, physical security may be an issue – in remote areas the sensors themselves may be susceptible to tampering. More importantly, the data may need protection against deliberate or accidental alteration. In some cases public access to the information from the sensors may be desirable and security mechanisms should not hamper this. Striking a balance between security and accessibility helps ensure that all parties can trust the systems.

The SPINS (Security Protocols for Sensor Networks) project (Perrig et al. 2001) developed a suite of security protocols optimised for sensor networks which included features such as data confidentiality, two-party data authentication, evidence of data freshness and authenticated broadcast for severely resource-constrained environments.

5 Exemplar ISNs

These three examples, drawn from environmental and well-being monitoring, illustrate complementary aspects of ISNs. Directly analogous to the DIID context, these technologies provide exciting new opportunities for monitoring the natural environment, in this case measurement of water levels and air pollution, and in monitoring movements of animals and humans. Sensor networks make it possible to deploy more devices in order to obtain more data more often, and this greater richness of data is set to create a powerful impact on environmental monitoring and decision making. Traditional solutions involve dataloggers from which data is collected periodically in person or via telemetry. With wireless communications and energy drawn from local sources such as solar cells, devices can be deployed without the constraints of having to wire them up or make them accessible, and data can be conveyed when needed.

5.1 FloodNet

Flood damage represents a major ongoing cost and risk may be increasing due to land-use change, climate change and flood-prone investment. When a flood occurs, the cost of damage has a clear correlation with both the depth of the flooding and the advance time at which warning is given. By applying pervasive computing technologies on the floodplain we have the potential to obtain better data from which to make predictions, and we can do this in a timely manner in order to improve warning times. Such systems provide an excellent illustration of the benefits and challenges of pervasive computing in the environment. Deployment is facilitated by wireless technologies but we have issues of power for the devices and the need for very long unattended periods of operation. This scenario also emphasises the energy optimisation challenge, because the data is most important during flood conditions and this is when solar energy is typically least available.

FloodNet (De Roure 2005) is a flood warning system that uses a set of sensor nodes to collect readings of water level and it communicates these via an asynchronous reliable messaging infrastructure to a grid-based flood predictor model. The reporting frequency of sensor nodes is influenced by local conditions and also the flood predictor model. This system is notable both for the adaptive sampling regime and the methodology adopted in the design of the adaptive behaviour, which involved the development of simulation tools and very close collaboration with environmental experts.

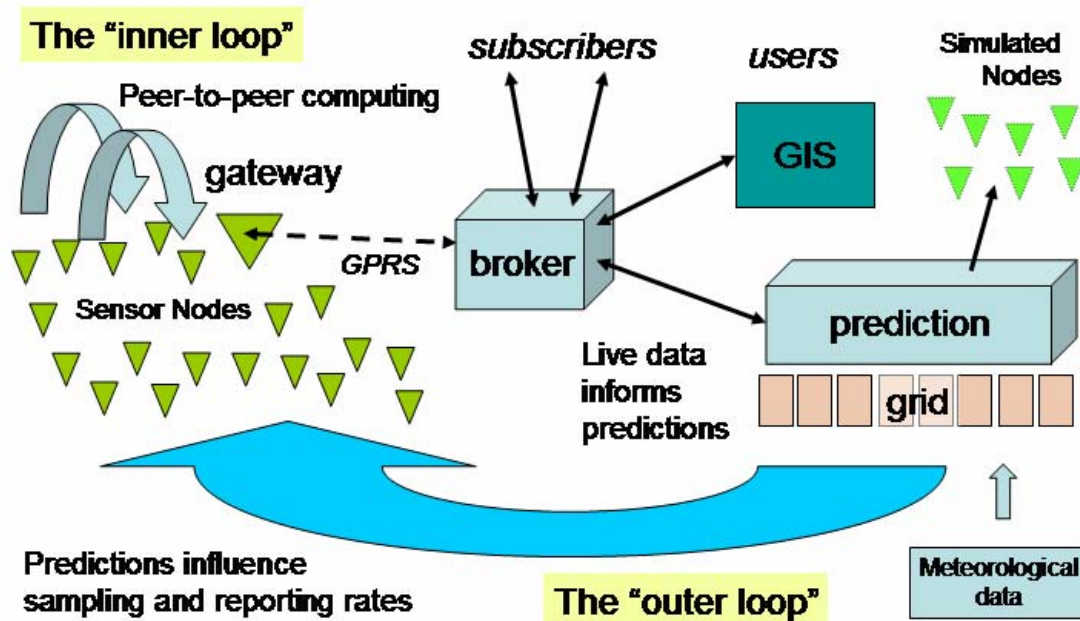


Figure 2: The inner and outer loop model

The adaptive sampling regime has an 'inner loop' and 'outer loop' as can be seen in Figure 2. The sensor nodes are able to communicate with their peers, agreeing such things as sampling and reporting rates. Partially processed data is then communicated to the broker. Once the data is outside the sensor network, additional processing can be carried out – in this example using a simulator on the grid as well as additional meteorological data. The predictions are fed back to the nodes in the sensor network (the outer loop) to provide additional information on which to base the sampling and reporting strategies. The use of the broker technology, developed for FloodNet by IBM, has proven very effective as it enables multiple users to subscribe to the appropriate incoming data streams.

FloodNet is investigating what is possible with relatively powerful sensor nodes and explicitly encoded intelligent behaviour. It is complemented by a sister project, Self-Organising Collegiate Sensor (SECOAS) Networks, in which the behaviour of the sensor network adopts techniques inspired from biology (Wokoma et al. 2005). SECOAS is exploring the future of low-power, low-cost, disposable nodes. In practice, it is anticipated that future sensor networks will involve a combination of these approaches.

5.2 Pollution monitoring

High concentrations or extended periods of exposure urban pollution can have serious health effects. Many local councils monitor urban pollution using a small number of sparsely distributed pollution sensors. These measurements

might represent a sample several square kilometres away, but it is known that pollution levels can vary from street to street. The Urban Pollution Monitoring project (Steed et al. 2003) has investigated local variations in pollution by using mobile sensors to detect the pollutant, carbon monoxide (CO). Carbon monoxide levels vary greatly depending on the local configuration of buildings and other environmental factors such as wind speed and direction. This additional metadata needs to be captured as part of the information about the network.

The approach attempts to bridge the gap between large, sparse sensor networks and small, dense studies. Data is collected from tracked mobile sensors that are carried by pedestrians or mounted on vehicles. These mobile sensors are not as well calibrated as fixed sensors, but they do give a spatial picture of pollution variation. The ultimate goal of this research is to produce a public infrastructure that can be shared or instantiated by any users wanting to investigate pollution. This would expand the system to many more pollution sensors, leading perhaps to more of a public understanding activity where people can build their own sensors and contribute data.

5.3 Well-being monitoring

Work at BT on well-being monitoring of the elderly (Hine et al. 2005) provides an excellent example of the use of sensors to monitor the well-being of individuals remotely. Sensors placed around the homes of elderly people living alone allow health and social care professionals to monitor their well-being and track changes in their patterns of behaviour. The data collected is used to build up a picture of daily activity which can be used to measure a person's ability to live independently. This also has the potential to prevent dangerous and debilitating incidents from occurring through timely intervention. However, monitoring behaviour in this way raises important ethical issues concerning consent and confidentiality that, although beyond the scope of this report, need to be considered.

In this application, and related applications tracking animals, we see the same pattern of data collection, transmission and processing, with the potential for intelligent behaviour – e.g. local analysis and processing – to prioritise response.

6 Lifecycle of ISNs

According to our holistic view, we need to look at the entire lifecycle of an ISN in order to fully understand their role in DIID.

6.1 Device design

There are often cost benefits to adopting off-the-shelf devices and adding hardware extensions using standard interfaces in order to provide specialised capabilities, rather than designing bespoke devices. For example, recent deployments of hand-held solutions have made effective use of mobile phones and PDAs. In contrast to legacy devices, newly emerging devices

tend to be highly reconfigurable by downloading new configurations and new software.

We anticipate new generations of devices with highly specialised capability, such as 'lab-on-a-chip' techniques to detect very low chemical concentrations. As technology evolves further, it is clear that devices will be smaller, and this enables new forms of deployment. Power is a significant issue, for example, with significant consumption due to many forms of data transmission, and is particularly critical in certain deployments; e.g. where devices are not accessed for considerable periods of time.

Design methods are becoming established to assist the creation of devices for human users, and there is increasing experience in designing devices for use in hostile environments. The research community is increasingly acknowledging the importance of hardware–software co-design.

6.2 Deployment

Typically, there is an existing deployment of legacy devices, and their associated processes, which may be coupled into a newly deployed system, manually or automatically – in practice, the heterogeneity of sensor networks is inevitable. Deployments usually feature redundancy, i.e. sufficient data can still be obtained should some devices fail. The security of devices, and authenticity of data, are important issues. As devices have become more sophisticated, multiple sensors may well be involved in obtaining information about the context of a sensor reading in order to assist with its interpretation.

There are many stakeholders in deployed ISNs. For example, there may be multiple sets of sensors running over the same infrastructure, and multiple applications consuming the data.

6.3 Configuration and maintenance

This phase must not be neglected and is essential to establishing dependability of the devices. It involves the ongoing test, calibration (where appropriate) and configuration of deployed devices. Devices should have a self-diagnostic capability and can be designed to self-configure according to their circumstances, an important capability where manual configuration of every device is not a desirable option or remote configuration is possible but fails – this is an aspect of 'intelligent' behaviour. When circumstances allow, a regular status report by devices facilitates automatic monitoring and automated response to failure.

6.4 Data processing and transmission

Devices may collect and transmit data at different moments. Some signal processing may be required. If a device is preconfigured to prioritise the data it receives from a sensor it may, according to priority, choose to store, compress, or take immediate action which gives rise to urgent transmission. Not only the detected data but also the associated metadata, which is needed

to interpret the data, must be transmitted. Some data analysis, processing or filtering may occur on the device.

A wide variety of communication mechanisms are now available. Broadly, they are 'store-and-forward' like email or 'synchronous' like a phone call. Only in the case of synchronous devices does it know immediately that data has been received. Networks and their associated infrastructure are increasingly able to make 'intelligent' decisions about routing content to appropriate points (content-based routing).

6.5 Information processing

There is a significant task of combining information from multiple sources, for purposes of decision support but also to make automated decisions to update device behaviours according to prevailing circumstances. For example, appropriate devices can be reconfigured to prioritise certain information according to the larger picture. This decision making is complex as incoming data quality may be variable and some readings will be false. It is important to track the quality of all data.

One approach to this function is to maintain a model (computer simulation) that is coupled to the real data and makes predictions about possible futures that assist decision making and also the automated configuration. This type of time-critical computation can be assisted by new infrastructures such as grid computing, especially appropriate with the increasing number and density of readings provided by sensor networks. Such simulations can also draw on historical data.

6.6 Notification and information

Pervasive computing technologies also have a role in alerting and informing users, with a degree of intrusiveness appropriate to the situation. Where users are required to follow a specific procedure in response to the alert, instructions can also be provided via the pervasive devices.

6.7 Information re-use

The information collected creates a digital record which is available for re-use. While it is possible to plan for anticipated re-use, there is significant benefit in unanticipated re-use. This reinforces the importance of capturing appropriate metadata, adopting re-usable representations, and thinking also about the issues of curation.

7 Summary

Sensors acquire information from the natural world. ISNs are smart devices and systems with additional processing functions beyond gathering data and transmitting it. They have capabilities for learning, adaptation and self-configuration, they are reconfigurable and they can perform data interpretation, correlation, fusion and validation. To achieve this, they typically

contain some limited processing functionality and may behave with a degree of autonomy. Data processing and, in particular, data transmission make demands on the power sources for the sensors, which is a significant issue in some deployments.

7.1 Key summary points

DIID context. ISNs are the front end of a pipeline of information processing, providing prioritised real-time monitoring which feeds into interconnected computer systems that process and manage the data. These systems respond to queries and conduct data analysis and mining that leads to knowledge about infectious diseases, provide decision support and facilitate emergency response. The detection and monitoring may involve new sensor devices and may also be integrated with other pervasive computing technologies and infrastructure such as mobile phones and intelligent buildings.

Information focus. The purpose of an intelligent sensor network is to produce information of known reliability and fidelity. Erroneous data from sensors can lead to unpredictable behaviour beyond the failure of individual devices, hence validation is important. Intelligent sensors can compensate for sensor errors, detect corrupted data, perform self-testing and respond to changing conditions. Metadata is just as important and ultimately supports search, retrieval and analysis, and increasing automation. There are issues of privacy and confidentiality both in the raw data and what can be inferred from it.

Intelligence and adaptive behaviour. There is an important interplay between the autonomic (self-managing) behaviour of the ISN and the way it is coupled into prediction models and decision support systems. For example, sample rates can vary dynamically according to local circumstances (sensor readings, available power) and may also be influenced by external requirements. Current systems do not exhibit a great degree of autonomic behaviour.

Communications. Current methods for transmission include wireless communications, such as shorter-range techniques (Bluetooth, zigbee), wireless ethernet and deployed wider-area infrastructure such as 3G. Within the DIID context, there is particular focus on prioritisation – what information can be stored and forwarded later, what needs to be sent immediately. There are also passive techniques – for example, information can be collected from a passive device using induction loops. Data transmission has significant impact on the use of energy sources and there is considerable research on energy-conserving techniques.

Current practice. There are examples of existing deployments. In the natural environment, we see non-intrusive and non-disruptive monitoring of sensitive wildlife and habitats, while in the built environment they include air pollution monitoring. Many of these are isolated, relatively small-scale deployments that do not need to integrate with other deployments. There are also military applications. The real deployments are essential for informing the research on

the individual components and also in enabling the data fusion and integration to be explored.

7.2 Future trends

Systems on chip and novel devices. Advances in silicon technology are enabling entire systems to be fabricated on a single chip (systems-on-chip), a transition from the traditional Application Specific Integrated Circuits (ASIC). Low-power design reduces power consumption to prolong battery life. ASICs are now being developed for smart sensor applications with built-in self-test and auto-calibration. Work on thick-film sensors is leading to many interesting devices, including self-powered microsystems (traditional batteries may not always be appropriate, especially in remote systems) and MicroElectroMechanical Systems (MEMS).

Scale. Technology trends in producing smaller and lower-cost sensors, and commercial drivers for pervasive computing devices, will inevitably lead to an increasing number and variety of deployed devices. Data will be obtained and processed from multiple sources. There will be increasing redundancy of information. Simulation techniques will be increasingly important in the investigation of large-scale systems issues during the design phase.

Interoperability. Discovering devices and their capabilities, composing these to perform useful tasks, and generating information for both anticipated and unanticipated re-use, all require interoperability. This will be achieved through standards for describing devices, services and data. This movement is in evidence through 'semantic web services'.

Autonomy. Devices will be increasingly self-configuring, self-calibrating, adaptive and reconfigurable. The mathematical and software techniques for this are being established through a large body of research in the multi-agent systems community. Autonomy is necessary to handle the scale of deployment and the increasing complexity of these systems.

Trust. Techniques for security, dependability, trust and provenance in pervasive computing devices are crucial in the DIID area. They are beginning to attract attention in sensor networks and will become more sophisticated. As well as being able to increase trust, the significant step is in knowing explicitly what level of reliance to apply to information provided by the network.

Usability. Considerable research effort is going into human-computer interaction issues with new forms of pervasive computing device. In addition to the devices themselves, usability issues apply to access to information from sensor networks and to the design, deployment, maintenance, debugging, re-use of the sensor networks themselves.

'In the wild'. With increasing numbers of deployed devices, new deployments can both interfere with and benefit from existing deployments. Deployment is an ongoing process and the deployed devices together form a heterogeneous distributed computer which is continuously changing. As a legacy of devices

builds up, there will be an increasing focus on re-use and re-purposing of deployed assets.

Acknowledgements

FloodNet is part of the Centre for Pervasive Computing in the Environment (EnviSense), funded by the Department of Trade and Industry (DTI) under the NextWave Technologies and Markets programme. The Urban Pollution Monitoring project is part of the EQUATOR Interdisciplinary Research Collaboration, funded by EPSRC (Grant GR/N15986/01). The authors thank Professor Tom Rodden, University of Nottingham, and Professor Ian Marshall, University of Kent, for their assistance in producing this review.

References

- Berners-Lee, T., J. Hendler, et al. (2001). 'The semantic web.' *Scientific American*.
- De Roure, D. (2005). 'FloodNet: a new flood warning system.' *Ingenia* 23.
- De Roure, D., N.R. Jennings et al. (2005). 'The semantic grid: past, present, and future.' *Proceedings of the IEEE* 93(3), 669–681.
- Estrin, D., R. Govindan et al. (1999). *Next century challenges: scalable coordination in sensor networks*. Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99), Seattle, Washington.
- Foster, I., N. R. Jennings et al. (2004). *Brain meets brawn: why grid and agents need each other*. Autonomous Agents and Multi-Agent Systems.
- Foster, I., C. Kesselman, et al. (2001). 'The Anatomy of the Grid: Enabling Scalable Virtual Organizations.' *International J. Supercomputer Applications* 15(3).
- Gao, W., L.M. Ni et al. (2004). *BLOSSOMS: A CAS/HKUST joint project to build lightweight optimized sensor systems on a massive scale*. Network and Parallel Computing, IFIP International Conference, Wuhan, China.
- Glynn-Jones, P., M.J. Tudor et al. (2004). 'An electromagnetic, vibration-powered generator for intelligent sensor systems.' *Sensors and Actuators* 110(1–3), 344–349.
- Hightower, J., S. Consolvo et al. (2005). 'Learning and recognizing the places we go.' *UbiComp 2005: Ubiquitous Computing: 7th International Conference, UbiComp 2005, Tokyo, Japan, September 11–14, 2005. Proceedings* 3660, 159.
- Hine, N.A., A. Judson et al. (2005). *Modelling the behaviour of elderly people as a means of monitoring well-being*. User Modeling.
- Intanagonwiwat, C., R. Govindan et al. (2000). *Directed diffusion: a scalable and robust communication paradigm for sensor networks*. MobiCom '00, Boston, MA.
- Laursen, W. (2006). 'Managing the mega flock.' *IEE Review*.
- Martinez, K., J. Hart et al. (2004). 'Environmental sensor networks.' *IEEE Computer* 37(8), 50–56.
- Min, R., M. Bhardwaj et al. (2002). 'Energy-centric enabling technologies for wireless sensor networks.' *IEEE Wireless Communications* 9(4), 28–39.

NGG3. (2006). *Future for European grids: GRIDs and service-oriented knowledge utilities: vision and research directions 2010 and beyond*. EU Grid Technologies Unit.

Perrig, A., R. Szewczyk et al. (2001). *SPINS: Security protocols for sensor networks*. Proceedings of Seventh Annual International Conference on Mobile Computing and Networks MOBICOM 2001, Rome, Italy.

Shah, R.C. and J.M. Rabaey (2002). *Energy aware routing for low energy ad hoc sensor networks*. Proc. IEEE Wireless Communications and Networking Conference (WCNC), Orlando, FL.

Shen, C., C. Srisathapornphat et al. (2001). 'Sensor information networking architecture and applications.' *IEEE personal communication*, 52–59.

Sinha, A. and A. Chandrakasan (2001). 'Dynamic power management in wireless sensor networks.' *IEEE Design Test Comp.*

Smarr, L.L., A.A. Chien et al. (2003). 'The OptIPuter.' *Communications of the ACM* 46(11), 58–67.

Steed, A., S. Spinello et al. (2003). *E-science in the streets: urban pollution monitoring*. eScience All Hands Meeting.

Teillet, P.M., A.E. Dudelzak et al. (2001). *A framework for in-situ sensor measurement assimilation in remote sensing*. Proceedings of the 23rd Canadian Symposium on Remote Sensing, Québec City, Québec.

Weiser, M. (1991). 'The computer for the 21st century.' *Scientific American* 265(3), 94–104.

Wokoma, I., L. Shum et al. (2005). *A biologically-inspired clustering algorithm dependent on spatial data on sensor networks*. 2nd European Workshop on Wireless Sensor Networks (EWSN-2005).

Wong, S.C., S. Miles et al. (2005). 'Provenance-based validation of e-science experiments.' *Proceedings of 4th International Semantic Web Conference (ISWC'05)* 3729, 801–815.

All the reports and papers produced within the Foresight project 'Infectious Diseases: preparing for the future,' may be downloaded from the Foresight website (www.foresight.gov.uk). Requests for hard copies may also be made through this website.

First published April 2006. Department of Trade and Industry. www.dti.gov.uk

© Crown copyright