Deploying Wireless Sensor Networking Technology in a Rescue Team Context

Ben McCarthy[1], Socrates Varakliotis[2], Christopher Edwards[1], and Utz Roedig[1]

[1] Computing Department, Infolab 21, Lancaster University,
London, WC1E 6BT, UK
{b.mccarthy,ce,u.roedig}@comp.lancs.ac.uk
[2] Computer Science Department, University College London,
Lancaster, LA1 4WA, UK
s.varakliotis@cs.ucl.ac.uk

Abstract. The computing department at Lancaster University are currently involved in the ongoing deployment of an advanced communications system designed to support the requirements of search and rescue teams. This system is based around the concept of using an all IP infrastructure to provide multi-functional data communications (such as group voice calling, live video streaming and location updates) to highly mobile vehicles and personnel in challenging environments. In addition to these types of data communications there is also a requirement to reliably transmit different types of sensor data information from the individual rescue team members, their vehicles and the casualties they locate and rescue. In this paper we describe the work we have carried out to incorporate an IP based sensor networking approach into our existing communications system deployment that we have in place with the Morecambe Bay Search and Rescue Team, in order to support Mobile Sensor Networks. In addition, we present results from our experimentation with our deployment that is specifically focused on the issue of wireless interference that our Mobile Sensor Networking solution is potentially subjected to.

1 Introduction

In general, search and rescue services have an inherent reliance on their communications solutions because their success is often closely related to their ability to accurately exchange information about a rescue situation between their team members. However, most search and rescue teams are still vastly under provisioned with regards to their communications equipment, with most teams still relying purely on push-to-talk radio solutions. Lancaster University are currently involved in an on-going effort to develop a powerful new communications solution that leverages Internet technologies to provide search and rescue teams with advanced functionality such as real-time localisation and mapping, real-time video streaming from vehicles and individuals, advanced voice services (i.e. tailored group calling, rather than indeterminate broadcasting of voice calls) and

P.J. Marron et al. (Eds.): REALWSN 2010, LNCS 6511, pp. 37–48, 2010.

real-time delivery of sensor data from the field of operation. In this paper we specifically focus on the experiences we have gained from incorporating an all IP sensor networking solution into our current mobile communications deployment with the Morecambe Bay Search and Rescue team (UK) [1]. In particular we provide an overview of the design of our solution as a whole and its reliability for delivering timely health statistics in challenging environments about individual rescue team members and the casualties they locate, as well as environmental data recorded by rescue team vehicles.

Our efforts with the Morecambe Bay Search and Rescue team demonstrate a real world deployment of Mobile Sensor Networking (MSN) technology and illustrate how sensor networks deployed in this context may often have to be considered as a component of a bigger overall system that has the potential to cause interference to their ability to deliver timely sensor data. From our deployment experience we have encountered scenarios where radio communication in over-lapping frequencies has caused sensor data delivery to be affected and in a critical scenario such as search and rescue the loss of reliable health statistics about a specific team member could be extremely damaging. The issue of radio interference arises because of our system's joint reliance on multiple wireless communication technologies that each transmit in the 2.4GHz ISM band of radio frequencies, namely the 802.11g, Bluetooth and 802.15.4 protocols. To investigate this problem further we constructed a number of tests designed to identify the real effects on sensor data delivery that are experienced with our system when used in the rescue team's operational environment.

The rest of this paper is presented as follows: In Section 2 we detail the nature of the rescue services deployment we are currently involved with and highlight how our system architecture maps onto their operational model. In Section 3 we focus on the way we have provisioned for sensor networking in the mobile context the rescue team operate in. In Section 4 we outline the specific hardware we used, as well as present the in-field interference experimentation we carried out. In Section 5 we provide an overview of related work that has been carried out in this field. Finally, in Section 6 we conclude with a discussion about our findings and overall experience of incorporating sensor networking into our search and rescue system deployment.

2 Trial Deployment

The Morecambe Bay Search and Rescue Team are a non-profit organisation that operate search and rescue missions primarily on and around the Morecambe Bay area in the North-West of the UK. In total their team consists of 16 operational members, all of whom are voluntary workers, devoting their time and effort for free. The Morecambe Bay area is the largest expanse of inter-tidal mudflats and sand in the United Kingdom and in total covers an area of 310 km^2. At low tide the sea retreats leaving a massive expanse of open sand that people, vehicles and animals use to walk/travel on and even work on (picking shellfish). In this state the bay area is deceptive and appears safe to access, but in reality it is peppered with treacherous pockets of quicksand and hidden channels that quickly fill areas with sea water when the tide returns. Due to these conditions many unsuspecting people (sometimes with vehicles) and animals get trapped, at which point their lives are immediately in extreme danger. For this reason the Morecambe Bay Search and Rescue team was created, to be able to cope with the specific requirements of rescuing people in the bay area's adverse and dangerous environment. More recently, the Morecambe Bay Search and Rescue team now also provide official support services to the area's fire brigade in land based situations which require their specialist expertise and equipment.

2.1 Rescue System Deployment

The rescue team currently have four primary vehicles that they use in their rescue operations (shown in Figure 1): 2 Hagglund BV206 All Terrain Emergency Rescue Vehicles, 1 Landrover 4x4 ambulance and 1 high-speed airboat. These vehicles carry out search operations (in incidents where a casualty's location is not already known), transport the rescue team members to an incident area and also act as a mobile command post for operation coordination. The communications system we have deployed is based around the notion of interconnecting wireless Vehicle Area Networks projected in and around each of the rescue team vehicles with wireless Personal Area Networks projected by each of the individual team members. Into each vehicle we have fitted a ruggedized, bespoke, vehicle specification Mobile Router (MR) which is powered from the vehicle's main power supply but that also trickle-charges a dedicated battery pack for use when the vehicle's own power supply is cut off. Each vehicle mounted MR has a number of wireless interfaces (each fitted with a dedicated external roof mounted antenna) including 2 802.11b/g interfaces (1 for interconnecting with other MRs, 1 for projecting a connectivity hotspot around the vehicle) and a cellular data modem capable of connecting to GPRS, EDGE and HSDPA services. In addition, we have also fitted the airboat vehicle with an INMARSAT BGAN vehicle terminal which is capable of maintaining a consistent connection to INMARSAT's broadband data service while the vehicle is in motion.

Fig. 1. Morecambe Bay Search and Rescue Team Vehicles

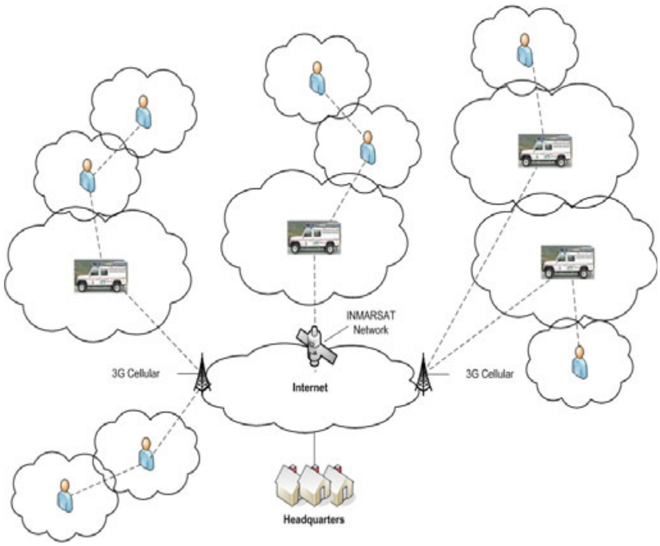

Fig. 2. Communications Model

In our model, as well as each vehicle containing a MR, each individual team member also carries a similar device designed in a suitable form factor for personal use. In particular, apart from being physically smaller than the vehicle mounted MRs the personal MRs also have a completely self-contained rechargeable Li-ion power supply and an additional 802.15.4 wireless interface. Once operational all of the MRs attempt to establish their own direct connection to the Internet via their cellular interface whilst simultaneously forming connections with other MRs via 802.11g. This process results in an interconnected Mobile Ad hoc Network (MANET) like the one illustrated in Figure 2 where any nodes capable of communicating directly with each other do so via multi-hop 802.11 and all other communication is performed using the Internet as a backbone. Using the Unified MANEMO Architecture (UMA) routing approach (discussed in further detail in the following section) allows us to provide unchanging, globally reachable IPv6 addresses to every node in the network, i.e. both the MRs and any host devices that subsequently connected to them, irrespective of how the topology of the network changes. This in turn allows data to be transmitted between every member of the rescue team, whether they are located at the headquarters, in one of the vehicles or out in the field, permitting communication such as group voice calls, streaming video delivery, location plotting and real-time sensor data monitoring. In the following section we focus in on how we have achieved this real-time sensor data delivery by incorporating Wireless Sensor Networking (WSN) technologies into our system to produce a fully functioning Mobile Sensor Networking (MSN) deployment.

3 Rescue System Sensor Networking

The term "Mobile Sensor Network" (MSN) is loosely used to describe WSN solutions which are capable of moving freely, potentially utilising different points of attachment to the Internet to continuously transmit real-time data about an entity or its surrounding environment as it changes location. This type of network of sensors is therefore ideally suited to gathering a broad range of data about a specific mobile entity, such as a vehicle or person, allowing their status to be continually monitored irrespective of their movement. In essence what we have achieved by incorporating our sensor networking solution into our search and rescue team deployment is a fully integrated real-world example of vehicle and person based MSNs. From a communications perspective, our approach achieves this by integrating the use of two key network layer technologies. Firstly the Unified MANEMO Architecture (UMA) [2] is used to provide the inter-communication capabilities of the rescue system, and therefore ensures that any IPv6 communication can be transmitted to any remote data sink in the Internet. Secondly, the 6LoWPAN adaptation protocol [3] is used to permit the efficient transfer of IPv6 packets to and from the resource constrained sensor nodes that are connected to the vehicle/person area networks via their low power 802.15.4 links.

The common features of WSNs are low bandwidth, constrained memory and limited computational power. Initially manufacturers introduced proprietary protocols to drive WSNs with customised link-layer solutions, assuming that IP was too resource-intensive to be scaled down to operate on the micro-controllers and low-power wireless links used in WSN settings. The 6LoWPAN protocol has addressed this situation and is what we use to provide the individual sensor nodes with IPv6 connectivity[3]. With 6LoWPAN, packet transfer from a sensor node to the network via a gateway is achieved by first fragmenting large IPv6 packets into chunks of 127 bytes or less. Once all fragments reach the gateway, packet re-assembly takes place and the composed IPv6 packet is subsequently routed to the Internet. The most commonly used header fields of the original IP packet may also be compressed as they are not required for routing within the sensor network, if layer 2 meshing is used. This compression and header stacking along with cross-layer optimisations result in low overheads, which translate to efficient transmission of IPv6 datagrams over low power networks. The overall savings can reduce the complete standard IPv6 packet (40 byte headers) down to an optimised few bytes only (around 2 bytes, at best, in typical uses) for Wireless Sensor Networks.

Fundamentally, UMA is a technique designed to enable mobile networks to perform persistent, uninterrupted IPv6 communication over the Internet, regardless of their potentially changing location and Internet access connection. In addition, UMA also ensures that mobile networks of devices can interconnect and communicate directly or share their Internet connections with other networks that cannot obtain their own Internet access connection, thus proliferating the availability of Internet access over a greater area. UMA achieves this by employing a technique that leverages the global connectivity characteristics of the NEtwork MObility Basic Support (NEMO BS) protocol [4] with the localised multihop communication support provided by MANET protocols. With UMA, every Mobile Router is registered with a corresponding Home Agent (HA) that records the its changing point of attachment to the Internet. This HA is located in the home network of the Mobile Router (i.e. the location it originates from, such as rescue team's headquarters) and intercepts packets destined for the mobile network whenever it is not directly attached to the home network. As the Mobile Router moves and potentially roams across different access networks or in-directly utilises another Mobile Router's Internet connection, it updates the HA with its new attachment point to the Internet. The HA then forwards all packets destined for the Mobile Router's network (either directly to it or in-directly via any other Mobile Router that is providing it with an Internet connection) via a bi-directional tunnel. This approach therefore keeps the mobility of the network transparent to any nodes it communicates with other than the HA and also prevents the traffic it generates from being Ingress Filtered in the access networks it visits. In total, this ensures that packets sent to and from the mobile network can use the same persistent IP address range to communicate regardless of its underlying mobility, and provides a highly robust solution because redundant, heterogeneous links to the Internet can be established and utilised if/when existing links fail.

4 Hardware/Software Setup and Experimentation

In one of our previous studies [5] we presented results from our initial lab based experimentation that acted as a proof of concept for our MSN solution and demonstrated the early potential of our approach. In this paper we confirmed the successful implementation of our solution using a Lippert Embedded Systems [6] CoolMoteMaster device, augmented to operate as a UMA Mobile Router also. This device however was not suitable for actual deployment in our scenario due to its form factor (too bulky and power hungry for personal use) and its non-ruggedized design. So instead we set about incorporating the concept we had proven in a lab environment into our existing ruggedized Mobile Router devices that we had already designed for real world deployment. To achieve this we needed to successfully incorporate an 802.15.4 interface into our Mobile Router design that could also perform the role of a 6LoWPAN gateway node and handle the packets sent too and from the sensor nodes appropriately. Both the gateway node and the sensor nodes are based on the popular Tmote Sky device (Telos Rev.B hardware platform, or 'mote'), the gateway node is directly attached to the Mobile Router board and powered via USB, whereas the sensor nodes are powered by battery. In each case we run the Contiki open-source operating system [7] which features the uIP network stack for IP communication with the motes. The uIP stack on the sensor nodes was further extended with a *sensor-side implementation* of the 6LoWPAN adaptation layer. On each sensor node there is a running measurement software component that upon initialisation performs basic IPv6 address auto-configuration by using a predefined network prefix known to the gateway. The measurement component then records information from the appropriate sensors, for the vehicle it records environment data

Fig. 3. SmartLife Technology "HealthVest"

including light, humidity and temperature measurements from the Tmote's embedded sensors. For the rescue team members we have interfaced the Tmote board to a prototype intelligent garment made by SmartLife Technology [8]. Their garment, known as the "HealthVest" (shown turned inside out in Figure 3) is a close fitting undergarment that incorporates their patented woven sensor technology to continuously monitor heart rate and electrocardiography (ECG) information.

In both cases, once this data is captured it is then placed in IPv6 packet payloads and transported with UDP to our remote sink application (a remote server located at the headquarters which provides the end user with combined mapping and personnel/vehicle localisation and sensor monitoring) in the following way. First, IPv6 packets larger than 127 bytes are appropriately fragmented by the 6LoWPAN sensor-side module on the source node for transport within the sensor network. The global IPv6 addresses in the headers are compressed as the sensor network only uses 16-bit link-layer identifiers. Once all related packet fragments reach the gateway, the corresponding gateway-side 6LoWPAN module re-assembles them into a full IPv6 packet and adds the decompressed destination IPv6 address of the target sink. The complete IPv6 packet is then passed to the kernel for further processing (i.e. IP routing using UMA, towards the data sink).

One final hardware device that is significant to our deployment and the experimentation presented in this paper is the Nonin Onyx II 9560 pulse oximeter [9]. This device is the world's first wireless enabled fingertip pulse oximeter and it works by periodically transmitting heart rate and blood oxygen level information over a paired Bluetooth connection. The ability to remotely monitor pulse oximetery data related to a casualty was a requirement specifically requested by the Morecambe Bay Search and Rescue team. This device offered us the perfect solution to this requirement but unfortunately operates over a different wireless technology to the others we already support and therefore its integration into the rescue system introduced further considerations about the wireless spectrum availability.

4.1 In-Field Experimentation

In total, at any one time in our rescue team communications system there can be up to 5 different wireless interface types operating at once. Whilst the satellite and GSM wide-area backhaul connections that we use in the system operate in distinct frequency ranges that do not cause interference with any of the other radios (INMARSAT BGAN receivers transmit and receive between 1525.0–1559.0 MHz and 1626.5–1660.5 MHz, O2 UK 3G service operates between 2125–2135 MHz), the 802.11g, Bluetooth and 802.15.4 interfaces all operate in the 2.4 GHz ISM band of frequencies. Each of these radio technologies transmitting in the 2.4GHz band have the potential to create contention with each other for access to the radio medium and therefore cause communications disruptions. This fact is further exacerbated by the proximity to each other that these interfaces can be expected to naturally operate in during a typical mission. To further elaborate on this statement we must first consider the primary purpose of these interfaces and then consider the context in which they are used. For example, in a typical rescue scenario, the 802.11g interface may be used extensively to carry multiple different types of communication (voice, video and location/sensor data). At any given time it may be relied upon to transmit relatively bandwidth intensive communication streams, therefore occupying one or two channels of the 802.11 radio spectrum which in turn can equate to potential interference across around 8 of the 802.15.4 channels. This type communication will also be relatively long lived in comparison to periodic sensor data transmission, and so it is extremely likely that sensor data transmission will occur at the same time as intensive 802.11 activity. One further communications consideration in our scenario occurs when a casualty is located and is transported back to safe ground and then hospital for treatment. As described in the previous section we have provisioned for the use of wireless pulse oximeters with our system during in-field operation that permit pulse rate and blood oxygen levels to be transmitted to awaiting hospital staff. To support this we auto-establish a serial-over-Bluetooth connection between the pulse oximeter and the onboard Bluetooth interface integrated in the rescue team member's Mobile Router. This means that from the moment the pulse oximeter is placed onto the casualty's finger, a third (periodic) transmission is introduced to the 2.4GHz ISM band. What this overview of the wireless communications landscape created by our rescue system shows is that it is therefore extremely important to ensure that radio interaction in our system operates fairly and that wireless transmissions do not unacceptably affect one another to the detriment of the entire system.

In order to test the interference effects experienced by the sensor networking data transmitted in our deployment we setup a series of experiments that involved the use of a number of different sensor networks and a number of different Mobile Routers in a range of different wireless configurations. In each test we deployed 1 Vehicle Mobile Router and then between 1 to 3 Personal Mobile Routers (introducing one after another to incrementally increase the level of interference experienced), with a simple hierarchical topology where all Personal Mobile Routers connected directly to the vehicle. We then connected a wireless 802.11 webcam to each of the personal area hotspots that each team member projected from their Personal Mobile Router. We then drove traffic over both wireless interfaces of the Personal Mobile Routers by streaming live video from the webcams up to the vehicle cabin. Finally, we activated the data sink in the vehicle also (in our deployment, all vehicles and the headquarters accept and display location and sensor data to the coordinators) and then configured each specific test. We configured the devices in each of our tests to forcibly create three different levels of combined interference: High, Medium and Low. These interference levels refer to our manipulation of channel selection for each of the 802.11g interfaces and the 802.15.4 interfaces, where:

- High: 2 802.11g on channel 7, 802.15.4 gateway/nodes on channel 18.
- Medium: 2 802.11g on channel 7, 802.15.4 gateway/nodes on channel 20.
- Low: 2 802.11g on channel 7, 802.15.4 gateway/nodes on channel 26.

In keeping with the focus of this paper on real deployment experience we decided to present our findings based on the most important aspect to the end user, i.e. the data arrival rate at the data sink. In our deployment (as with any other) the rescue team coordinators are interested in whether they are receiving sensor data correctly or whether they are not. As our sensor networking approach is fully IP compatible we opted to measure packet arrival rate using the Wireshark network protocol analyser which provides the ability to filter out specific packet types, offers accurate packet arrival times and also has additional tools for overall analysis. For each test, the sensors periodically transmitted their data once every second for a duration of 1000 seconds and we repeated the motions of the test 5 times over to gain average packet loss rates. Then at the end of each test run we connected a Bluetooth pulse oximeter to one of the rescue team member's Mobile Router and began transmitting casualty health statistics over each resulting network topology to observe whether the pulse oximetry data was affected.

The percentage loss experienced at the data sink in each of the configurations we tested can be seen in Figure 4. What is immediately evident from these results is the obvious correlation between the high level of sensor data loss and the high level of radio interference present. When operating on completely none-overlapping radio frequencies the level of packet loss experienced was encouragingly low. When operating in rescue missions, team members originating from one vehicle often work together in a designated area. Each vehicle will tend to carry around 4 rescue team members, 3 that will alight at given locations and 1 driver that remains in the vehicle at most times. With this operational model it would be straightforward to devise a suitable channel separation scheme between each of the team members belonging to one vehicle that would ensure their radio communications were not in over-lapping frequencies for the bulk of their time in the field of operation. Less encouraging however were the levels of loss we recorded at the "medium" level of interference. In a planned frequency allocation model like the aforementioned channel separation scheme based on vehicle assignment, this configuration effectively represents the encroachment of radio transmissions from devices outside of the planned model, i.e. when other vehicles converge in a single location. This scenario is not atypical and must therefore be expected to occur, in which case our results show the level of successful delivery of sensor data really begins to suffer.

Interference Level:	1 Mobile Router	2 Mobile Routers	3 Mobile Routers
Low:	1.4%	1.3%	1.8%
Medium:	4.6%	16.7%	27.1%
High:	42.6%	44.5%	46.0%

Fig. 4. Percentage loss in each configuration

One interesting result was the negligible effect that adding additional radio interference to the already high level of over-lapping 802.11 transmissions had when further Mobile Routers were introduced. The radio interference effectively appeared to reach a saturation point that it never rose significantly above. One explanation for this behaviour could be the specific webcam hardware used in our deployment. At present we use Panasonic webcams (because of their native IPv6 support) these webcams stream their video captures using TCP and as a result will aggressively back off under heavy interference conditions. Away from the average loss rates entirely, we also observed that during the "High" level of interference sensor data packet loss was rarely consecutive. When a sensor data packet was lost, in most cases the following packet in the transmission (i.e. the next 1 second interval) could be expected to arrive. This was observed with noticeable consistency to the point where in tests that yielded almost 50% loss, packets could be seen to arrive in an almost a 1 on, 1 off fashion.

5 Related Work

At present we are unaware of any successful deployment attempts that have aimed to provide IP based Mobile Sensor Networking in a challenging environment such as search and rescue. However, whilst work related to our system deployment as a whole is scarce at the moment, there is a lot of existing research related to the radio interference experienced between 802.11 and 802.15.4 which is a subject we focused on in this paper. One of the closest studies to the analysis we carried out in our deployment was a recent paper specifically focusing on the interference properties experienced between two sensor nodes communicating with one another in a Body Area Network (BAN) [10]. The authors of this paper provide a very complete and thorough analysis of the effects of interference on a point-to-point 802.15.4 link when it encounters 802.11 traffic in a similar radio frequency. They are able to draw stark parallels between an increased successful transmission rate and increases in the transmit power that the sensor nodes operate at.

Additional work that is related to the particular sensor data that we transmit in our deployment has also been carried out from the perspective of medical applications in hospitals [11]. In this study the authors set out to determine the suitability of 802.15.4 for the purpose of carrying health statistics information in a typical hospital environment. However whilst the sensor data and end devices discussed in this study is related, the operational environment is naturally very different and the authors attained their results through simulation.

As well as studies analysing the impact of 802.11 on 802.15.4, there also exists studies that have considered this problem from the opposite angle of how

802.15.4 can impact on 802.11's overall capability. Interestingly there are papers with conflicting conclusions in this research space [12] [13] with studies finding no effect and another recent study finding effects on 802.11 in specific circumstances. In our deployment any effect on 802.11 by 802.15.4 was not noticeable and will unlikely ever be if we continue to operate a model of 1 second sampling times and single sensor gateway/node pairings between each individual Mobile Router. For the foreseeable future it seems that this model is sufficient for our requirements and at that node density the 802.15.4 transmissions simply are not frequent enough to cause significant disruption to any 802.11 communications.

6 Conclusion

In this paper we have provided an overview of our mobile networking deployment activities with the Morecambe Bay Search and Rescue Team. We have focused in particular on our recent work related to incorporating an innovative all-IP sensor networking technique into the existing IP infrastructure established by our rescue system communications solution. Apart from detailing the technologies and hardware/software components we have used to achieve these outcomes we have also provided an insight into some of the experiences we have gathered so far. Building on these experiences we have also outlined some in-field testing we performed to ascertain the suitability of our approach for reliably transmitting important sensor data information throughout the rescue team network. Specifically, this testing concentrated on the potential negative effects of radio interference that can be experienced in large, multi-function communication systems like the one we have deployed.

In particular, we found radio interference from intensively used 802.11g wireless interfaces did have a negative impact if channel overlap existed with the 802.15.4 spectrum that was utilised. However, what we also observed was that even in very high levels of interference there would often be little consecutive loss. Therefore when a packet of sensor data was lost, in all but the worst cases the next packet could statistically be expected to be delivered. This means that the importance of the data being carried must be taken into consideration, and more specifically, the criticality of every single reading produced. If we take our deployment for example, the significance of every single sensor data reading is debatable. Certainly from the perspective of an environmental reading from a vehicle the levels of loss experienced are unimportant, but this can even be said of the health data of the individual rescue workers. During an operation a team coordinator will monitor the progress of their team looking at location and movement, if the weather is bad or worsens they may casually maintain a watch on the overall health statistics of their members, but any data loss would most probably go unnoticed. However, if a team member was identified as going overboard into deep water in Winter this interaction could quickly change and every reading could become extremely important. Related to this is the criticality of a casualty's health statistic data, however whilst this data is seen as being extremely important the use of Bluetooth and its Adaptive Frequency Hopping Algorithm in this case suitably addresses this problem.

Finally in this paper we showed that the effects of radio interference between these technologies have been researched before, however not in a mobile communications system as comprehensive as our deployment. This has additional significance because in our system we can potentially control each of the sources of interference and attempt to minimise its negative effects through collaborative channel selection. Using our understanding of the interference problem domain we can now attempt to develop a channel selection scheme that attempts to identify radio spectrum in the 2.4GHz ISM band that is already saturated and then dynamically adapt its use. This concept in itself is potentially very challenging, bringing additional pitfalls because of the highly mobile nature of our deployment, but if solved correctly could provide an effective tool for other mobile networking solutions of this nature as and when they start to be deployed.

References

1. Morecambe Bay Search and Rescue, `http://baysearchandrescue.org.uk/`
2. McCarthy, B., Edwards, C., Dunmore, M.: Using NEMO to Support the Global Reachability of MANET Nodes. In: Proceedings of Twenty Eighth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2009), Rio de Janeiro, Brasil, April 19–24 (2009)
3. Montenegro, G., Kushalnagar, N., Hui, J., Culler, D.: Transmission of IPv6 Packets over IEEE 802.15.4 Networks. IETF RFC 4944 (September 2007)
4. Thubert, P., et al.: NEMO Basic Support Protocol. IETF RFC 3963 (January 2005)
5. McCarthy, B., Edwards, C., Varkliotis, S., Kirstein, P.: Incorporating mobile sensor networks into the internet - an all ip approach. In: Proceedings of the 4th International Workshop on Mobility in the Evolving Internet Architecture (MOBIARCH 2010). ACM, Chicago (2010)
6. Lippert embedded systems homepage, `http://www.lippertembedded.de/index.php`
7. The Contiki operating system, `http://www.sics.se/contiki`
8. Smartlife technology homepage, `http://www.smartlifetech.com/`
9. Nonin mdical inc, onyx ii 9560 wireless pulse oximeter homepage, `http://nonin.com/PulseOximetry/Fingertip/Onyx9560`
10. Handziski, V., Hauer, J., Wolisz, A.: Experimental study of the impact of wlan interference on ieee 802.15.4 body area networks. HotEmNets (2010)
11. Cypher, D., Golmie, N., Rebala, O.: Performance analysis of low rate wireless technologies for medical applications. Computer Communications 28(10), 1255–1275 (2005)
12. Howitt, I., Gutierrez, J.: Ieee 802.15.4 low rate - wireless personal area network coexistence issues. In: Wireless Communications and Networking, WCNC 2003, vol. 3, pp. 1481–1486. IEEE, Los Alamitos (2003)
13. Pollin, S., Hodge, B., Tan, I., Chun, C., Bahai, A.: Harmful coexistence between 802.15.4 and 802.11: A measurement-based study. In: 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications, CrownCom 2008, pp. 1–6 (May 2008)