Networked Appliances for Home Healthcare and Lifestyle Management

Paul Fergus, David Llewellyn-Jones, Madjid Merabti, Arshad Haroon

Networked Appliances Laboratory
School of Computing and Mathematical Sciences
Liverpool John Moores University,
Byrom Street, Liverpool L3 3AF, UK.
Email: {P.Fergus, D.Llewellyn-Jones, M.Merabti, Arshad Haroon}@ljmu.ac.uk;
A.Haroon@2006.ljmu.ac.uk

Abstract. Advances in technology and the increased use of home medical devices, will revolutionise the way public healthcare is administered. Homes and their associated networks in conjunction with such devices will take over many mundane healthcare tasks and manage new and enriched lifestyle choices that affect our overall quality of life. Through the combination of wireless and fixed networking infrastructures explicit links will form between the home and its devices and medical installations, such as hospitals. Through these interconnected networks new real-time healthcare management systems will emerge that continually provide information and react to adverse or unusual medical conditions received from occupants within the home. Achieving this will undoubtedly require a convergence between home networks and medical installations in order to harness the power afforded by both. We present a new approach using a working prototype to implement networked medical devices within the home capable of monitoring data received from an individual, which could then be accessed within the home and medical installations.

1. Introduction

Home networking has received considerable attention over the past decade with many research projects and blue chip organisations having tried to transform passive interaction within our homes to include interactive technology that helps attend to our every need. No longer is it acceptable to carefully configure home environments where such configurations restrict how and when we use technology. We have become more demanding, were we expect to use our technological solutions wherever we are. Moving from one room to the next, our preferences are expected to follow us, communications have to be ubiquitous, and multimedia access and control, as well as information services must be readily available. We are less tolerant towards the mundane interactions associated with home management. We want lights to automatically switch on as we enter a room and heating systems that self-adapt to environmental and physiological changes.

As such home environments have become superimposed with technology where intelligence is beginning to emerge to optimise the operation of such environments. This has made it possible to create complex configurations so that devices are networked and only operated given certain conditions. The health industry is beginning to ask questions about how such advances could be used to push secondary healthcare services into the patient’s home. Many non-networked devices have been developed such as blood glucose meters, blood pressure monitors and heart monitors that can be used easily by patients within the home. Such solutions have prompted further investigation into whether providing such devices with networking capabilities could significantly reduce costs. Necessitated by the fact that such networked devices exist, the hope is that large overlay networks will emerge that provide ubiquitous healthcare and lifestyle monitoring services similar to that illustrated in Figure 1.

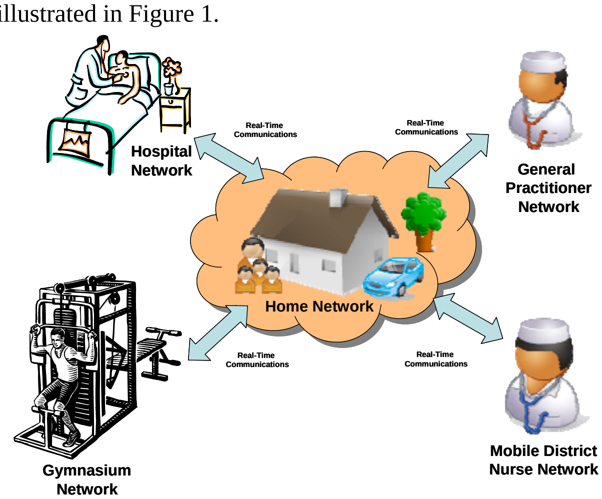

Fig. 1 Home, Healthcare, and Lifestyle Platform

Implementing such a solution could have many benefits. Perhaps it will help reduce unnecessary fatalities that could have been avoided if such technological solutions were available. In this sense networked medical devices will act to prevent adverse affects to home users by reacting to the detection of early symptoms. For example, many elderly people are not always aware that they are dehydrated. This can pose significant problems causing secondary conditions to develop, resulting in potential loss of life. For example, early detection could help reduce stress hormones called cortisol, which can suppress the production of white blood cells, which in turn may leave the patient susceptible to allergies. Here a very simple networked medical device could monitor hydration levels and prompt the person to take a drink. If the condition persists a healthcare practitioner could be notified via the home network.

Imagine an extension to this where lifestyle choices are carefully monitored and used to observe the short and long-term affects of such choices. On entering the gymnasium, medical devices automatically connect to the gymnasium network and begin collecting different biological and health related data, such as heart rate recovery information and the number of calories lost during a session. This could all be directly linked into personal profiles managed at your home, within your local hospital and by your general practitioner. Results from your recent gym sessions could determine what food you could have, with certain choices resulting in approval or a recommendation that alternative foods should be used. Tagging items with the use of sensor technology such as RFID make the implementation of such a system possible. The challenge is to overcome the many difficulties that currently exist within both the home networking and healthcare domains, such as interoperability and Food and Drug Administration (FDA) approval.

Addressing this challenge, we provide a platform that allows consumer devices and home medical devices to be more easily combined to meet the unique medical needs of the user and the environment. The proposed framework ensures users are not burdened with the details of this process; it simply happens unbeknown to them and without affecting their day-to-day home activities. For example, a user simply introduces new medical devices into the home and they automatically and wirelessly communicate and configure themselves with all other devices in the home without human intervention. We achieve this using a peer-to-peer network that consumes devices and the functions they provide, which themselves are abstracted as discoverable services. Using the peer-to-peer network these services can be automatically composed into any number of healthcare applications to administer healthcare solutions. In this way the integration and management of healthcare and lifestyle choices is achieved using these devices and services.

2. Background and Related Work

The home has been the focus of many research initiatives for several years and a great deal can be learnt from the existing work in this area.

More recently, industrial and academic efforts have given us many solutions resulting in the introduction of a wide spectrum of wired and wireless infrastructures and network protocols such as LonWorks, CEBus, SmartHouse, VHN, HomePHA, HomePnP, IEEE1394 (Firewire), X-10, IrDA, IEEE802.11b, IEEE802.15.4, Bluetooth and HyperLAN/2 [1]. However, despite the long list of advantages each provides, several challenges still remain to be addressed, most notably interoperability [2, 3] and the difficulties associated with the integration of combined functionalities, including the use of vocabularies to discover and describe devices and their operational capacity.

Lots of existing frameworks have been developed that aim to provide interoperability in the home environment such as the Digital Living Network Alliance

(DLNA) [4], the Open Services Gateway Initiative (OSGi) [5], Reconfigurable Ubiquitous Networked Embedded Systems (RUNES) [6], AMIDEN [7], and Universal Plug and Play [8]. None of these have specifically addressed the ad-hoc, functional capabilities that we believe will be of particular benefit in the medical domain.

Whilst many challenges still exist, enough advancement has been achieved within the home networking domain were the deployment of networked medical devices and services is now possible. Medical devices themselves will become increasingly network-enabled and support better capabilities. This will allow devices ranging from sensors to high-end multimedia appliances to form part of these networks, where the functions they provide can be pervasively distributed and used to create new and innovative home healthcare solutions. In support of this Lee *et al.* [9] make some interesting observations. They argue that realising such a complex integration will result in a ubiquitous heterogeneous overlay containing different protocols, different contexts, location and property independent device control, ad hoc device registration and updating, real-time data processing and alarm management, and high volumes of event traffic. Many of these issues have been the focus of much home networking research, consequently much can be learnt from these advances when integrated networked medical devices within the home.

Many have foreseen this gap in the market and are already building on the notion that the environment will become more finely augmented with information-bearing devices, such as sensors. For example, Stankovic *et al.* [10] believe that wireless sensor networks will play a big part in home healthcare provisioning. In what they refer to as in-home assistance and smart medical homes, this technology will help the aging population by providing memory enhancement; home appliance control; data collection; and emergency communication services. Wearable sensors will serve to collect data, which could be used by services and mined as a data source for next-generation clinical trials.

Sensor networking research is clearly a very active topic in both the home and medicine, with estimates already reaching $18 billion on BioMEMS research in 2005 as reported by Gupta *et al.* [11]. They argue that investment is significantly influenced by a drive to revolutionise healthcare to allow pervasive and continuous monitoring of real-time physiological conditions. There are many motivating factors for the use of biosensors. For example, sensors have already been designed for internal patient examination as is the case with SmartPill designed to monitor and transmit gastro-based information to a home computer. Furthermore, wireless sensors have been designed to collect routine data, such as Wireless ECGs, which would otherwise require trained medical practitioners. Using sensor technology to obtain biological data in this way is seen as a considerable cost saving to the medical organisation and at the same time an advanced means of processing and reacting to real-time data. Nonetheless, Gupta *et al.* report that implementation is plagued by many technical challenges. Specifically: heterogeneity between devices and sensors, interoperability between the numerous wireless protocols available, reliability, context, and real-time adaptation.

Despite the long list of challenges, we cannot ignore the potential benefits, and the fact that a new bread of networked medical device has great potential in moving from reactive to preventative medicine. Current practice means we tend to only visit a doctor when something has already gone wrong. Only at this point are technological advances used to diagnose – and if possible – rectify a problem. This is unacceptable given that many fatalities or serious illnesses could be prevented through early detection. For example, diabetes, if not treated, can result in the amputation of limbs. This could be addressed through the use of sensor technologies as discussed above using real-time pressure readings measured continuously and used to notify people that perhaps they need to put their feet up to avoid unnecessary risk to their health [12]. There are obviously many other preventative steps that could better utilise technology and significantly reduce health related risks. For example, Voss describes bandages that warn you when infection is detected, even identifying the bacterium and applying the appropriate antibiotic required to treat it.

This section has discussed two distinct research domains that are showing signs of convergence. Each approach in its own right provides benefits, but also suffers from drawbacks. The challenge is to capture the positive aspects from each to enable a new research strand allowing home networks to safely underpin the integration of networked home and medical devices.

3. Home Healthcare and Lifestyle Management Framework

As we have seen, many standards exist in home networking to provide a means of abstracting device functionality as services. Healthcare is beginning to turn to this domain in an attempt to capitalise on its benefits to offer home healthcare services. For example, Kin *et al.* have used OSGi as their deployment platform [13]. Coupled with the ability to create this abstraction, peer-to-peer technologies such as JXTA [14] have been used to provide a vehicle to dynamically disperse and discover healthcare services [15, 16]. Given the success of such approaches it is appropriate to build on these advances. Using networked appliances abstracted as peer-to-peer services we present an investigation detailing how medical devices can seamlessly interoperate within and across different networks.

The discussion describes how device operations are abstracted as services and deployed within the network, where interconnected medical devices are free to utilise these services. Given that all functionality is viewed as a collection of these services, compositions are automatically created based on shared characteristics. For example, a service (implemented on a device) that outputs multimedia streams may be composed with services that consume multimedia data. Given such compositions exist, mechanisms responsible for creating relationships can manage their execution allowing medical devices to self-adapt to unforeseen changes that occur within and between services and the environment. Combining these principles allows high-level applications to be automatically created through device and service interactions. This reduces the difficulties associated with combining medical devices making it easier for healthcare specialists and home users alike to simply use the devices without having to define how applications are created beforehand.

3.1 Approach Overview

The design goals provide the system requirements for a suitable scheme as described in this paper. The principle goals are as follows.

- Peer-to-Peer. Home and medical devices and services are published and utilised using peer-to-peer technologies.
- Abstraction. Device functionality is abstracted as services, which can be discovered and utilised by other devices or services within the network.
- Service Architecture. Devices have the ability to offer zero or more framework services. If a service is not hosted by the device then it discovers and uses the required service remotely within the network.
- Ad hoc Services. Home and medical devices offer and discover services without third-party intervention. Once devices are switched on they can offer their services without having to register them with a third-party registry.
- Semantic Descriptions. Services are described and discovered using semantic annotations. Semantic interoperability between different ways of describing services is achieved using an ontology (a shared understanding of some subject) [17].
- Automatic Device and Service Composition. Devices and services are automatically composed using semantic annotations and service capability models, i.e. devices that provide the same functionality may redundantly co-exist, and consequently it may be desirable to select devices that support the best configuration.
- Device and Service Management. Devices and services self-adapt and extend the functionality they provide beyond that which they were initially designed to do. Conflicts are detected and rectified as and when they occur.

The approach presented in this paper can be seen as an important step towards networked medical device research where the line of enquiry is likely to prove fruitful. The remainder of this paper explains how the principal goals have been incorporated within the framework and highlights the novelty of our approach.

3.2 Device Interoperation

The proposed scheme provides a middleware that operates between medical devices and services contained within different networks and the different applications created that use such devices and services. Figure 2 illustrates how all communication between compositions passes through the service integration framework and the technology adapters used to communicate with devices. Not all appliances are network-enabled. Consequently, technology adapters may be required to integrate legacy devices. Nonetheless, networked-enabled appliances may choose to bypass these adapters and interact with the framework directly.

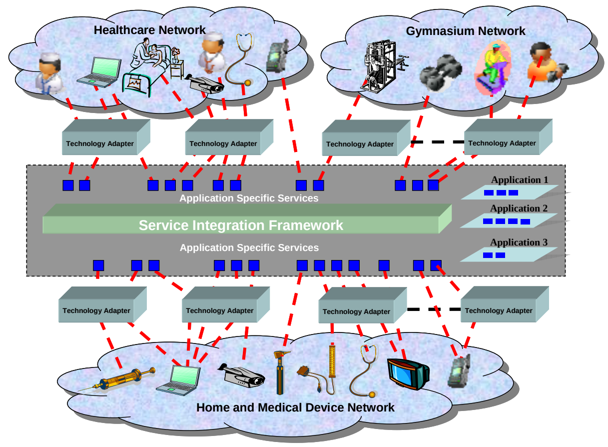

Fig. 2 Proposed Framework

As devices connect to the network they perform four tasks. First, they either publish the framework services they implement and/or discover and use framework services provided by other devices. Devices with varied capabilities must be accommodated. In this sense the framework needs to be flexible enough to allow devices to implement some, all, or none of the framework services. For example, a personal computer might implement the entire collection of framework services because it is capable of doing so, whereas a biofeedback sensor might not due to its limited capabilities. In the latter instance the sensor may discover and use framework services remotely (for example, services provided by a personal computer). Second, devices publish their operational capabilities (e.g. a blood pressure device may abstract the functionality it provides as services). Third, devices form relationships with each other and services within the network. This is not to be confused with step one which describes how a device implements framework services; this step relates to how devices are connected together. Fourth, devices self-adapt to connections between each other and services, as well as environmental changes.

Many home networking standards utilise message-based approaches to dynamically discover resources. Adopting a simplified viewpoint, discovery approaches such as the simple service discovery protocol [18], as used in UPnP, provide an oversimplified means of discovering content based on keyword matching. The problems associated with this approach are well documented, resulting in the use of semantics to try and address these limitations. Consequently, the approach taken in this paper extends message-oriented techniques to include the use of semantics to describe services and device capabilities.

The rules governing semantics differ significantly over keyword matching where logical constructs interconnect semantic elements. This does have problems of its own where heterogeneity exists between the ways concepts are defined. Given this limitation, semantic information does however provide a conceptual net in which a larger number of queries can be handled that would otherwise be missed. Whilst promoting the unconstrained use of different vocabularies, there is a need to provide semantic interoperability mechanisms to help resolve conflicts. Techniques used by the Semantic Web [19] and the Semantic Web Services communities can help to address this issue. We can therefore use unconstrained vocabularies in line with semantic interoperability mechanisms to help interconnect devices and services. In achieving this, the approach adopted by Paouluci et al [20], is extended to not only discover services, but also to help automatically compose services with little or no human intervention [21]. However, the precise nature of medical terminology means that particular care must be taken in this area to avoid inappropriate semantic matching.

Building on peer-to-peer technologies, the extension of such approaches to incorporate the use of semantics and self-adaptation enriches them to better enable devices to automatically interoperate. Using semantically annotated service advertisements devices can advertise the functionality they provide by publishing them within peer-to-peer networks. Depending on the technology, advertisements are likely to have a time-to-live value that expires after a given time period – consequently, advertisements have to be periodically re-published to keep the service alive. Devices may fail unexpectedly. In this case, service advertisements would no longer be re-published and would eventually be purged from the network. Early results suggest that our approach, based on these principles, helps improve medical device and service deployment as well as provide additional benefits that allow devices to self-adapt.

Devices can join the network as specialised or as simple networked medical devices. Specialised devices may have the ability to host services, store and evolve semantic information used to describe and discover services, as well as propagate service requests within the network. A simple networked appliance by definition may not have these capabilities. This type of device would typically join the network, propagate queries and invoke discovered services, or simply just be controlled. An example of such devices might be a sensor or an X-10 enabled lamp [22]. This would enable any device, irrespective of its capabilities, to effectively choose how it wants to interact within the network.

The Primary service is the only compulsory service each medical device is required to implement. This service is responsible for marshalling all communications between framework services (Secondary Services) and Application Peer Services (APS), whether they reside locally or remotely within the network. Simple networked medical devices, such as a Doppler tester, classed as a sensor (has very limited capabilities), could use the secondary services provided by a more general device (e.g. a personal computer). The Primary service on the sensor marshals communications to the hardware of the sensor and the Primary service located on specialised devices. The Secondary services themselves provide framework functionality, whilst Application Peer Services provide abstractions to the operational functions devices provide such as readings provided by the Doppler tester.

When devices initially join the network the goal is to try and discover devices they may have relationships with. For example, a treadmill machine in a gymnasium may discover all the equipment used by someone during a training session. This configuration may also include a medical interface used to submit information to the user's general healthcare practitioner and their mobile phone could be used to synchronise data readings with healthcare services on the user's home PC used to monitor fitness levels.

By managing references to external services (relationships), devices can adapt to changes that occur. For example, if the user's mobile phone becomes unavailable for some reason, the treadmill could detect this change and automatically select an alternative device, such as an Internet information processing service, provided by another device or the gymnasium. Looking at another example, we are able to monitor the osmolality of bodily fluids, which may be obtained using blood serum, urine or saliva sensors, to inform a person if they are dehydrated. This could be achieved via any audio device in the person's immediate vicinity. This has the potential to provide a breakthrough in all aspects of healthcare and lifestyle management where early detection can minimise the effects dehydration may have, which can often result in headaches and cramps [23]. Obviously the difference here is the ability to allow the sensors themselves to form relationships with other home devices and decipher how best to inform the user as and when required using appropriate services. In particular, care must be taken during the composition of services to ensure that the most appropriate functionality is chosen amongst ostensibly similar devices or services. In this instance compositions between devices and services need to be based on the best resources available. "Best" could mean resources that provide the best functionality, or resources that are free to use. Whatever the definition, the term needs to be identified and adhered to. In this paper, it refers to the service and the ability of the device to provide that service.

As a result, it is beneficial for devices to describe their capabilities in terms of how well they execute the services they provide. This means that devices would benefit from providing a corresponding device capability model describing how capable it is. By processing these capability models, a device can determine other devices' suitability before including them in a composition. For example, a mobile phone may be connected to the gymnasium's satellite system allowing a user to listen to a music channel. If the user moves to a different location a plasma display may be selected because it provides a better entertainment experience (audio and video). If the user moves location (to alternative gymnasium equipment) where the plasma display is no longer visible, the configuration may adapt to re-direct the satellite signal back to the mobile phone.

4. Secondary Framework Services

We turn our attention now to describing how the Secondary services extend the Primary service functionality discussed in the previous section to allow devices to be dynamically composed. The process requires the automatic resolution of terminology differences between vocabularies, and the assessment of device capabilities. Importantly the framework is not limited to these services but has the ability to accommodate any new functionality that may be required as the framework evolves.

4.1 Dynamically Composing Devices

To address the issue of increasing complexity, it is important to ensure that management tasks are considered in parallel with device development. We believe the management of devices and services to be perhaps one of the most important aspects of medical device research. For this reason, devices must be equipped with management techniques to automatically integrate and adapt services used in high-level applications. For example, when a networked-enabled dehydration device is connected to a home network, it must automatically advertise its services and form relationships with notification devices, such as those that provide audio capabilities, so that the person can be notified when signs of dehydration are detected. Underpinning this dispersed operational functionality, management services need to react to adverse environmental conditions and composite conflicts, such as the unavailability of services [21]. This requires a collaborative 'intelligence' allowing devices to negotiate these interactions and composite conflicts.

One approach is to enable machine-to-machine negotiation using service ontologies [24], to semantically describe functionality so that such functionality can be discovered and composed to create high-level tasks or applications (service compositions). Using the concept of Inputs, Outputs, Preconditions, and Effects (IOPEs) devices and services are composed by matching similarities between the service request and the service ontologies. Service ontologies, in conjunction with domain ontologies are used to match vocabularies that may be syntactically different but semantically equivalent. This approach overcomes many of the limitations associated with simple syntactic attribute-value pair matching.

Whilst using semantics in this way provides an interesting line of enquiry, it is well documented that automatically composing devices is problematic. This can be attributed to the variation in how service interfaces are defined and described, where one single difference in the parameters used in the signature can render the service inappropriate for the composition. To accommodate this, mechanisms could be adopted similar to constructors used in object-oriented programming where a base constructor could be used and then extended to include the different ways the object can be created. Whilst this is one possible solution, it is not scalable. A more effective way may be to extend the concept of automatic service composition to enable signatures to be composed, resulting in new signatures emerging. We achieve this using intermediary services and extended interfaces. For a more detailed description of this please refer to [21].

4.2 Enhancing the User Experience

To fully utilise devices we need to determine what functionality is available and how effectively devices providing those functions can execute them. For example, a plasma television and a video-enabled mobile phone are both able to process video information. However, the quality of these services will differ. A plasma television may be considered better equipped to process high definition video compared to a mobile phone. Given these different capabilities selecting and composing devices must be dependent on quality of service and application specific requirements. To address this, devices have the ability to select devices and services based on how well they can execute a service [21]. This can be achieved using metadata models, dynamically generated to assess device capabilities. Each model contains information about every resource the device manufacturer deems important. Using this information will allow a device to rank functionality based on ordered capabilities where the top of the list provides highly capable devices and the bottom those that are less capable. For an in-depth discussion of this algorithm the reader is referred to [21]

This section has considered all of the secondary services that comprise our proposed framework. The following section demonstrates their practical use in an intelligent home environment.

5. Implementation

Our framework has been designed to equip medical devices with plug and play capabilities, and aid interoperability between heterogeneous services they provide that can self-adapt to environmental changes. Although still at an early stage, we have built a prototype to evaluate the design, which demonstrates how state changes in the human body can be used to discover and control networked devices or alert a healthcare practitioner of the event. To achieve this we use an ECG sensor, a camera, a mobile device and a high-end visual display unit. The ECG sensor provides a service capable of taking ECG readings and making that data available within the network. The camera is used to stream images of the patient to either the mobile device or the visual display. The mobile device is carried by a medical practitioner, whilst the visual display is located in the practitioner's place of work.

When devices are initially switched on they perform three tasks. First, they publish the services they provide using the Primary Service discussed above. Second, they try to discover devices within the environment they have a relationship with (in the prototype the sensor service discovers the ECG service, the camera service, the mobile device and the visual display) using the Second Services to

compose service interfaces and determine device capabilities. Third, they monitor the communication links they have with other devices and using our framework devices self-adapt to changes that occur.
The prototype demonstrates how high-level applications can be altered at runtime when exceptions occur within compositions. For example, when the healthcare practitioner moves out of the immediate vicinity of the visual display, an event is triggered to inform the framework that this device should no longer be used as part of the composition, and an alternative visual display should be used, such as that provided by a mobile phone.

The Primary and Secondary services have been developed using the JXTA protocols. JXTA allows any device to be connected to the network independent of platform, programming language, or transport protocol. This includes simple medical devices such as our Doppler tester. Framework services (Primary and Secondary) are predetermined and each device understands how to discover and invoke them. Bindings between devices and Secondary services are achieved using pipe advertisements. Pipes are used by services and applications to send and receive information between the two. Through the abstraction provided by JXTA they can be viewed as virtual communication channels, much like sockets programming, except pipes exist independent of location and network topology.

Application-specific services are used to expose the operational capabilities of all devices. Many functions may exist, consequently pipe advertisements used are not known by devices beforehand. As discussed above, service ontologies describe device functionality by mapping high-level semantics onto low-level service interfaces. This is achieved using the service ontologies [24].

We used a BIOPAC MP150 data acquisition workstation [25] to read the electrical activity produced by the heart in real-time, which was received using electrode sensors (LEAD1102-R, LEAD110S-W and LEAD11), connected to a MEC110C extension cable [25]. The electrode sensors themselves are connected to a Biopotential amplifier (ECG100C Electrocardiogram Amplifier) and the human body using Vermed SilverRest Electrodes.

Communication between the MP150 and the heart monitoring service is achieved using a standard straight through RJ45 Ethernet cable. The heart monitoring service communicates with a TINI microcontroller [26] in ad hoc mode using a wireless 802.11b access point. The access point is directly connected to the TINI board using a straight through RJ45 cable, which simulates an on-board wireless interface, which to date is not implemented natively on TINI microcontroller boards. The microcontroller is connected to a PC via its serial port, where a service is used to turn the camera on or off and receive AV streams. This is achieved using the JXTA pipes and the Real-time Transmission Protocol. These streams are piped towards the visual display or the mobile device, dependent on where the medical practitioner is. For example, if on call and out on a visit, the mobile phone is used; if at work the visual display is used. The described configuration is illustrated in Figure 3.

In a situation where the medical practitioner is away from the visual display, the actions require that the screen resolution be adapted to support a mobile device, the media format set to 3gp and the content is transcoded from H.264 to 3gp. The configuration is dependent on where the medical practitioner resides and our framework can dynamically change which device the medical practitioner is using. The framework provides self-adaptation functions to detect such changes and reconfigure compositions to best suite the medical practitioner's context.

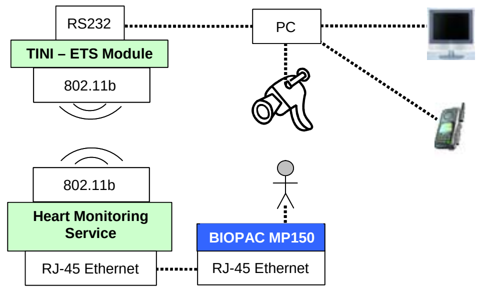

Fig. 3 Prototype Hardware Configuration

In the prototype the human body is connected to the BIOPAC MP150 via the heart monitoring electrode sensors. The values received from the sensors are analogue values between 0 and +10 volts. These are measured by the ECG 100C device and converted into digital values using the MP150. This digital data is transmitted to the Heart Monitoring Service, via the RJ45 connection and processed using two algorithms. The first algorithm analyses the digital data received from the BIOPAC and calculates the global maximum in the heart beat signal as illustrated in Figure 4. When processing the data we avoid local maximums by only processing peaks above 5 volts. The peaks in the data are determined by calculating the increasing values; when the value starts to decrease we know that we have reached the maximum point.

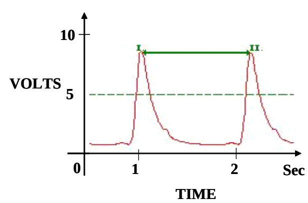

Fig. 4 Heat beats signals received from the ECG 100C

The second algorithm calculates the time between two peaks (also illustrated in Figure 4) and calculates the number of beats per minute (bpm). The bpm value is then used to determine if an alarm state has been reached. For experimental purposes, values between 50bpm and 75bpm are classed as normal heartbeat activity and values above or below this range indicate an alarm state has occurred. In this case a counter is invoked and incremented by one for each heartbeat out of range. If this counter reaches a predetermined value (for experimental purposes this value was set to 5), the Heart Monitoring Service generates an "enable-alarm-message" and sends it to the TINI microcontroller. If the heartbeats remain within range for 5 consecutive beats an “alarm-disable-message” is sent to the TINI microcontroller. Using a counter in this way enables the service to avoid triggering alarm conditions if spurious spikes in the data occur. For example it is possible for a spurious heartbeat to spike above 75bpm, therefore counting 5 consecutive beats ensures that an alarm condition is definitely present.

Messages are sent to the TINI microcontroller using 802.11b communications and are processed by an Ethernet to Serial (ETS) converter service written in Java, in order to control the underlying functions provided by the camera. In reality the TINI microcontroller and the camera would be one self-contained unit, however in the prototype the TINI microcontroller is connected to the camera via a PC. Upon receiving control messages from the Heart Monitoring Service the TINI microcontroller instructs the camera to activate. This results in the audio and video received from the camera being streamed to either the visual display or the mobile device.

Stankovic et al. [10] describe several application areas where wireless sensor networks could be used – sleep apnea, journal support, and cardiac health. We have successfully implemented within our framework mechanisms to obtain data through EEG and ECC sensors building on the ideas Stankovic et al. provide.

6. Conclusions

In recent years significant advances have been made independently in the areas of medical technology and networked appliance research. We believe that by combining elements from each of these domains there is the potential to provide further benefit, notably in relation to monitoring and preventative medicine.
We present a framework as an initial attempt to provide such capabilities. The framework has been tested using a prototype heart monitoring service able to stream video to a medical practitioner in the event that an alert is triggered, and which is able to react dynamically to real time changes in the situation and environment.

Whilst the medical devices and service framework used have each been proven effective in their own respective domains, this work represents a novel combination that provides additional benefits. We believe the direction represents a potentially crucial advance in medical technology that can – for example – build upon ongoing developments in home networking. However, although a working prototype has been achieved, we acknowledge that significant challenges remain, particularly in the areas of guaranteeing safety and satisfying regulatory requirements. We do not underestimate the importance of these, and aim to focus on these areas in our future work.

Acknowledgements

The authors would like to thank Lorna Bracegirdle from the NHS North Liverpool Primary Care Trust in the United Kingdom, for providing her valuable experience and understanding of medical device usage and healthcare provisioning within the home.

References

[1] B. Rose, Home Networks: A Standards Perspective. IEEE Communications Magazine, 2001. 39(12): p. 78-85.

[2] O. Abuelma'atti, M. Merabti, and B. Askwith, Internetworking the Wireless Domain. Proceedings of the 3rd International Symposium in Communication Systems, Networks and Digital Signal Processing (CSNDSP), 2002, Staffordshire, UK, p. 344 - 348.

[3] T. Zahariadis and K. Pramataris, Multimedia home networks: standards and interfaces. Computer Standards and Interfaces, 2002. 24(5): p. 425-235.

[4] DLNA: Overview and Vision. 2004, DLNA, Accessed: October 2006, http://www.dlna.org/about/DLNA_Overview.pdf.

[5] The OSGi Service Platform - Dynamic services for networked devices. 2005, OSGi Alliance, Accessed: 2006, http://www.osgi.org/.

[6] C. Koumpis, L. Hanna, M. Anderson, and M. Johansson, Wireless Industrial Control and Monitoring beyond Cable Replacement. 2nd Profibus International Conference, 2005, Warwickshire, UK.

[7] M. Minoh and T. Kamae, Networked Appliances and their Peer-to-Peer Architecture AMIDEN. IEEE Communications Magazine, 2001. 39(10): p. 80-84.

[8] UPnP Forum. 2005, Microsoft Corp., Accessed: 2006, http://www.upnp.org/.

[9] M. Lee and S. Kang, Multimedia Room Gateway for Integration and Management of Distributed Medical Devices Workshop on High Confidence Medical Device Software and Systems (HCMDSS), 2005, University of Pennsylvania, Philadelphia, PA, USA: University of Pennsylvania.

[10] J. A. Stankovic, Q. Cao, T. Doan, L. Fang, Z. He, R. Kiran, S. Lin, S. Son, R. Stoleru, and A. Wood, Wireless Sensor Networks for In-Home Healthcare: Potential and Challenge. Workshop on High Confidence Medical Device Software and Systems (HCMDSS), 2005, University of Pennsylvania, Philadelphia, PA, USA: University of Pennsylvania.

[11] S. K. S. Gupta and L. Schwiebert, Dependable Pervasive Health Monitoring. Workshop on High Confidence Medical Device Software and Systems (HCMDSS), 2005, University of Pennsylvania, Philadelphia, PA, USA: University of Pennsylvania.

[12] D. Voss, Smart Home Care: New diagnostic devices could save an ER visit. The MIT Technology Review, 2001: p. 31.

[13] N. Kin, Y. Jeong, S. Song, and D. Shin, Middleware Interoperability based Mobile Healthcare System. The 9th IEEE International Conference on Advanced Communication Technology, 2007, Phoenix Park, Gangwon-Do, Republic of Korea: IEEE Computer Society, p. 209 - 213.

[14] L. Gong, JXTA: A Network Programming Environment. IEEE Internet Computing, 2001. 5(3): p. 88-95.

[15] B. Lim, K. Choi, and D. Shin, A Secure Peer-to-Peer Group Collaboration Scheme for Healthcare System. 5th International Conference on Computational Science, 2005, Atlanta, GA, USA: Springer, p. 346 - 349.

[16] B. Lin, K. Choi, and D. Shin, A JXTA-based Architecture for Efficient and Adaptive Healthcare Services. International Conference on Convergence in Broadband and Mobile Networking, 2005, Jeju Island, Korea: Springer, p. 776 - 785.

[17] M. Uschold and M. Gruninger, Ontologies: Principles, Methods and Applications. The Knowleddge Engineering Review, 1996. 11(2): p. 93-155.

[18] Y. Y. Goland, T. Cai, P. Leach, J. Gu, and S. Albright, Simple Service Discovery Protocol/1.0, I. E. T. Force, Editor. 1999, Internet Engineering Task Force: Internet Engineering Task Force. p. 1 - 18.

[19] T. Berners-Lee, J. Hendler, and O. Lassila, The Semantic Web. Scientific America, 2001. 284(5): p. 34-43.

[20] M. Paolucci and K. Sycara, Autonomous Semantic Web Services. Internet Computing, 2003. 7(5): p. 34 - 41.

[21] M. Merabti, P. Fergus, O. Abuelma'atti, Y. Heather, and C. Judice, Managing Distributed Networked Appliances in Home Networks. To appear in the Proceedings of the IEEE, 2008.

[22] Z. Yuejun and W. Mingguang, Design of Wireless Remote Module in X-10 Intelligent Home. IEEE International Conference on Industrial Technology, 2005, Hong Kong: IEEE Computer Society, p. 1349 - 1353.

[23] R. L. Gunter, W. D. Delinger, T. L. Porter, R. Steward, and J. Reed, Hydration level monitoring using embedded piezoresistive microcantilever sensors. Medical Engineering & Physics, 2005. 27(3): p. 215 - 220.

[24] OWL-S 1.0 Release. 2003, DAML, Accessed: 2006, http://www.daml.org/services/owl-s/1.0/.

[25] BIOPAC Systems, Inc. 2004, BIOPAC Systems Inc., Accessed: 13 May 2004, http://www.biopac.com/.

[26] Introducting TINI: Tiny InterNet Interface. 2003, Dallas Semiconductor, Accessed: 18-11-2005, http://www.maxim-ic.com/.