

A Context Aware Wireless Body Area Network (BAN)

Tony O'Donovan¹, John O'Donoghue², Cormac Sreenan¹, David Sammon², Philip O'Reilly², Kieran A. O'Connor³

¹Mobile and Internet Systems Laboratory, Dept. of Computer Science, University College Cork, Ireland.

²Business Information Systems, University College Cork, Ireland.

³Mercy University Hospital, Cork, Ireland.

Abstract—In monitoring a patient's real-time vital signs through Body Area Networks (BAN), rich data sources are communicated to medical practitioners. The benefit of BANs may be negated if medical practitioners are overloaded with streams of BAN data. It is essential that data is delivered in a timely context aware manner. In this paper a BAN designed for falls assessment among elder patients (65+ years) is presented, with an emphasis on the communication scheme chosen. The FrameComm MAC protocol described in this paper employs three data management techniques, 1) message priority, 2) opportunistic aggregation and 3) an adaptive duty cycle, all of which are designed to ensure that patient vital signs (i.e. data packets) are delivered under a variety of network loads. The protocol is evaluated using a small laboratory network, initially configured to communicate Beat-to-Beat (continuous blood pressure) readings when a patient goes from a sitting to a standing position and then with added ECG (ElectroCardioGram) readings.

I. INTRODUCTION

Wireless sensor network based technologies can deliver a vast array of health-related services and large quantities of real-time patient centric medical information. A ubiquitous medical computing environment is where advanced wireless technologies (e.g. BANs) are applied to ensure the timely delivery of accurate patient based medical data [1]. As a BAN is capable of generating large quantities of real-time data, the potential for information overload is very high. It is imperative that only data which is of relevance to the medical practitioners is communicated in a timely manner. The challenge resides around two key areas 1) when does the BAN communicate data and 2) what data should be communicated. The layout of our BAN is given in Fig. 1. It contains two accelerometers (to identify the patient's current stance i.e. lying, sitting or standing), an ECG sensor (for cardiovascular activity) a Beat-to-beat sensor (for continuous blood pressure) and finally a pulse oximeter sensor (for pulse and blood oxygen levels).

In this paper, three BAN data management techniques are evaluated within the context of falls assessment among the elderly, they are 1) priority (i.e. in an emergence situation, which sensor readings or data packets get priority over other sensor readings), 2) aggregation (combining the data from several messages into one in order to reduce overhead and network contention and congestion, this allows more data travel through the network and helps ensure that the server or medical practitioner receives critical data) and 3) adaptive duty cycles (based on a specific set of circumstances the duty




Fig. 1. BAN Configuration.

cycle durations may be increased or decreased to meet the application requirements, thus preserving the life of the battery while meeting the needs of the medical application).

It is possible for a patient to experience a sudden drop in blood pressure when changing position, i.e. when going from a sitting to a standing position etc. This is a result of forcing the heart to pump more blood around the body to support the physical exertion of standing up. The patient can experience this as feeling faint but may ignore it. If such events continue but go undiagnosed then these fainting spells may worsen, resulting in a blackout and ultimately a fall.

A context aware system was implemented to evaluate the BAN with the above scenario in mind. The system has two network states, normal and high priority. The normal state represents the patient at rest, monitoring and logging readings from the sensors at a low data rate and duty cycle. Short periods in the high priority state are triggered when the patient changes position, (e.g. stands up, sits down), resulting in changes in data rate and duty cycle to allow more frequent sampling of sensors and increase available bandwidth respectively. By having such context awareness the network can operate in an energy efficient manner when the patient is at rest. However, it can also increase the sample rate and available

bandwidth when the patient stands up and a drop in blood pressure is most likely, allowing more data to be logged.

For the purposes of this paper the BAN, and FrameComm in particular, is evaluated initially with Beat-to-Beat (continuous blood pressure) readings only while in the high priority state. Readings from an ECG (ElectroCardioGram) sensor are included in a subsequent run.

The remainder of the paper structure is as follows; in Section II the motivation for the BAN and its data management techniques is given with particular reference to falls assessment. Related work is given in Section III highlighting current BAN trends and MAC protocols. Section IV presents the FrameComm protocol architecture with its data priority, data aggregation and adaptive duty cycles features. The evaluation of the FrameComm protocol is given in Section V under both configurations. Finally a conclusion on the effectiveness of the FrameComm protocol is outlined with a look as to where it may be applied in the future.

II. MOTIVATION

The demands of an aging population is one of the key drivers of the patient monitoring market in Europe. For these demands to be met, innovative method need to be developed. According to the European Commission a 60% increase is expected in the number of older patients by 2050; therefore aging is a big issue which will drive the patient monitoring market in Europe [2]. This aging population will have a tremendous effect on healthcare; particularly regarding care of the critically ill (more than half of all intensive care unit stays are incurred by patients over 65). “The integration of wireless technology into medical devices to capacitate easier information flow coupled with faster workflow and cost containment is likely to set pace for a future of health care that is highly patient-centric, prevention-based health care model” [3].

The eHealth industry in the European Union was worth close to €21 billion in 2006 and is expected to be the third largest industry in the healthcare sector, with rapid double-digit growth anticipated by the end of the decade. Also, strategies are published by the European Commission to accelerate the growth of the eHealth market in Europe. In fact, there has been a 60% reduction in acute care costs because of eHealth acceptance in Europe [2]. A fall is often defined as “inadvertently coming to rest on the ground, floor or other lower level, excluding intentional change in position to rest in furniture, wall or other objects” [4]. Falls and fractures are a major cause of disability, mortality and suffering for older people and their relatives, making fall prevention extremely important in countries with aging populations. Studies have identified that falls is one of the defining characteristics of older people and that 1 in 3 people over 65 and 1 in 2 people over 80 fall every year [4]. Older people aged 65+ in the Republic of Ireland currently account for approximately 11% of the total population and this proportion is expected to increase over the next twenty years [5]. Therefore, the economic and societal burden of falls is set to increase by

epidemic proportions in the coming decades. Currently, in the context of older people in the Republic of Ireland, there is an enormous cost to the health services (for example hospital and social care) with regard to falls and fractures. In fact, Irish research indicates that falls related injuries in older people cost the Irish economy €402 million per year [5].

If current trends continue costs will escalate, it is estimated that by 2030 the cost will be €1,500 - €2,000 million per year. It has been argued that a 15-30% reduction in falls can be achieved by implementation of a planned health and social care approach [4]. A number of medical conditions can greatly contribute to patient falling e.g. poor eyesight, heart arrhythmias or sudden blood pressure fluctuations brought on by exertion. A number of technologies exist which can alert care givers or medical practitioners of a patient’s state of health in particular among the elderly the ‘European Seniorwatch Observatory and Inventory’ examined the adoption and usefulness of BAN type technologies highlights the their risks and challenges [6].

A change in blood pressure is of particular interest to medical practitioners as it may lead to a dizzy spell or a fall. With this in mind the blood pressure sample rate is increased when a change is most likely to occur, i.e. after a patient stands up. When the patient is at rest a lower sample rate is sufficient. This context aware approach reduces power consumption and network congestion without sacrificing valuable information.

In the local context, when older people present themselves at the falls clinic as day patients or in-patients the Consultant Physicians in Geriatric Medicine can investigate changes in heart rate and blood pressure with changes in posture (within a controlled medical environment – i.e. wired devices). Using the current method, older people are strapped on a tilted table and are wired to devices that are used for monitoring changes in heart rate and blood pressure (Beat-to-Beat) while the position of the table is changed and visual monitoring of position is recorded. This examination lasts for a maximum of 40 minutes only, is time consuming to setup and labour intensive (generally requiring 2 medical practitioners to conduct the examination activity). Furthermore, the frequency of such examinations is limited to 8 to 10 a month. By introducing a BAN device to capture and access the patient vital signs directly after a patient fall patient examinations would be conducted more frequently (as it would be less time consuming and more closely related to real life human mobility or Activities of Daily Living (ADL)) which would provide a greater volume of real-time patient-specific data with new richness in terms of falls monitoring and assessment.

III. RELATED WORK

A correctly designed wireless BAN can provide medical practitioners with valuable real-time and historical information. It is important that a patient is completely comfortable and accepting of such technologies if they are to be used in everyday life [7]. Presented in [8] are a number of issues surrounding the usage of wireless body area networks 1) Human Factors “So Comfortable That the User Can Wear It for Long Periods” i.e. having a comfortable wireless sensor

network (WSN) that the user voluntarily wears throughout the working day enables the medical practitioners to capture vital sign readings which may not have been possible with other clunky more inhibiting devices and 2) “Light Weight, Low Power, Low Cost and Operate Under Various Environmental Conditions” i.e. a light weight WSN improves user comfort and allows for multiple sensors to be comfortably worn with little or no interference with ADL.

A BAN can generate large quantities of data over a period of time. This data may need to be offloaded to a second location for further processing. [9], [10] and [11] demonstrate the feasibility of transferring BAN data onto a more traditional network e.g. GPRS within medical settings. With such an approach data can be communicated to the relevant medical practitioners for further analysis. In [12] the DMS architecture was developed to tackle a number of the data management tasks which need to be executed at a number of stages i.e. managing the data from the BAN, to a central server to the medical practitioner. The FrameComm protocol, first described in [13], presented in this paper focuses on managing the BAN data in relation to the patient’s positional change (i.e. sitting to standing) and the effect such a change may have on the BAN’s ability to cope with low to high sensor sampling rates.

In [14] three techniques in relation to MAC (Medium Access Control) protocols were evaluated within a BAN context: 1) with S-MAC a scheduling approach is used to coordinate sleeping modes among neighbouring nodes to limit idle listening states [15]. 2) A preamble sampling technique known as LPL (Low Power Listening) is used in the WiseMAC protocol, where a node regularly samples/polls the nodes for a brief sample period to test whether a packet needs to be received [16]. When the node wakes up and senses no traffic, it can revert back to sleep mode. If the source has to send some data, it transmits a wake up preamble. Finally 3) an ultra-low duty cycle SCP-MAC is achieved by combining scheduling and polling techniques [17]. Very low duty cycles are achieved by synchronizing the polling times of all neighbouring nodes. Here scheduled Polling-MAC eliminates long preambles in low power listening.

IV. ARCHITECTURE

This section describes the components that make up the BAN.

A. Mote

The DSYS25 hardware platform is a Lego-like 25mm x 25mm stackable system. Its modular nature lends itself to the development of numerous layers for use in different application scenarios. Layers can be combined in an innovative plug and play fashion and include communication, processing, sensing and power. The communication layer is comprised of an ATmega128L micro-controller, a ChipCon CC2420 Zigbee compliant radio transceiver and a simple wire antenna. An FPGA for high-speed DSP forms the processing layer while various application specific sensors, as well as a generic sensor interface/communications layer, have been developed for the

sensing layer. The power layer may include batteries or other energy supply mechanisms such as solar cells.

The modular nature of the platform means a node can be tailored to a particular application by adding a suitable sensor layer and power layer. The nodes for the BAN will consist of the communication layer, a battery power layer along with appropriate sensors.

B. Sensors

Each node in the network fulfils a different role which will determine the sensors it will be equipped with.

1) *Motion Detection - Accelerometers*: The key feature of the application is its ability to detect a change in the patient’s position and alter sampling rates and network conditions accordingly. Since one of the triggers for a sudden drop in blood pressure is the patient going from a sitting to a standing position the detection of such a change is crucial. In order to detect this transition the BAN includes two nodes equipped with an ADXL202 [18] 2-axis accelerometer.

The ADXL202 is a small, low power sensor on a single chip with a direct interface for microcontrollers, all of which make it ideal for use in WSNs. It can detect $2mg$ motion with a range of $\pm 2g$ making it more than capable sensing the movement of a human body.

2) *Blood Pressure - Beat-to-Beat*: A continuous blood pressure monitor is employed to sense changes in a patient’s blood pressure. Portapres [19] is an advanced non-invasive blood pressure measurement and Beat-to-Beat haemodynamic monitoring system and was the sensor chosen for this. It can return finger arterial pressure and heart rate, reconstructed brachial arterial pressure and various cardiac parameters.

3) *ElectroCardioGram (ECG)*: An electrocardiograph (ECG) is one of the key indicators used by medical practitioners to assess the health of the human heart. An ECG is a very sensitive indicator of the health of the heart and is widely utilised as a diagnostic procedure. It is designed to record/monitor the activity of the human heart which produces tiny electrical pulses that spread through the heart. This electrical activity forces the heart into a continuous cycle of contraction and relaxation states thus pumping blood around the body to vital organs. The heart’s electrical impulses can be detected by an ECG device e.g. Holter or wireless Body Area Network (BAN). For our emulated experiments the ECG readings were modelled against the Tyndall-BAN [20] which is based on a 3-lead ECG layout 1) positive point, 2) negative point and 3) ground point (Fig 1).

C. TinyOS

TinyOS is an open source operating system for wireless sensor networks. It has a component based architecture that allows rapid development of applications and it minimises code size as required by the memory limitations of wireless sensor nodes. The standard distribution includes various components, such as MAC and routing protocols, and supports many commonly used platforms and sensors.

The BAN will use the TinyOS operating system with a customised network protocol.




Fig. 2. Framelet-rendezvous.

D. Network Protocol - FrameComm

There are several requirements to be considered when designing a network protocol for an application such as this. As with all wireless sensor network communication protocols it must be energy efficient and relatively straightforward, this is due to the hardware constraints inherent in WSNs along with the difficulty in debugging. In addition the scheme should support multiple message types, variable traffic loads and prioritisation.

FrameComm, originally described in [13], is a duty-cycled CSMA based MAC protocol that conforms to all these criteria. It is an implementation of the Framelet rendezvous mechanism extended to include message prioritisation, opportunistic aggregation and an adaptive duty cycle. Details of the Framelet concept and the features of FrameComm follow.

1) *Framelet Rendezvous*: Many low power communication protocols employ a duty cycle in order to reduce idle-listening and increase energy efficiency. Since a sensor node whose radio transceiver is in sleep mode cannot receive messages some method of establishing transmitter receiver rendezvous is required. One method of achieving rendezvous involves the retransmission of a series of identical packets, which we call *Framelets*, in such a way as to guarantee at least one framelet coincides with the receiver's listen activity. This scheme's simplicity in design makes it easier to implement and debug than many of the proposed protocols.

The basic mechanism can be seen in Fig. 2. The sender transmits a series of n framelets, where, δ , denotes the transmission time of a single framelet and, δ_0 , denotes the time between successive framelets. The receiver meanwhile alternates between sleep and listen states for a duration of Δ_0 and Δ respectively. To ensure rendezvous the following inequalities must hold:

$$\Delta \geq 2 \cdot \delta + \delta_0 \quad (1)$$

$$n \geq \left\lceil \frac{\Delta_0 + 2 \cdot \delta + \delta_0}{\delta + \delta_0} \right\rceil \quad (2)$$

The receiver can acknowledge receipt of a Framelet as shown in Fig. 2. This allows the sender to cease transmission of framelets once rendezvous is established. A more detailed description of Framelet rendezvous can be found in [21].

2) *Interrupt Concept*: All CSMA based communication protocols require a carrier sense mechanism, in order to avoid collisions, prior to any transmission a node must first check the channel to discover whether or not a neighbouring node is already sending. Since framelet transmission includes a gap between successive framelets (for acknowledgments) a simple Clear Channel Assessment (CCA) is insufficient. The CCA could occur during one of the gaps between framelets and would incorrectly indicate a clear channel. With this in mind FrameComm instead uses a *pre-send listen* as a carrier sense, that is it listens to the channel for a time of Δ , ensuring any ongoing transmission is received.

In addition to acting as a carrier sense mechanism this *pre-send listen* provides a node, say Node A, with information about the type of traffic currently occupying the channel (Node B). The header of any overheard message can be examined to obtain attributes such as its destination, payload length and message type. By comparing these attributes with those of the message awaiting transmission Node A can decide if it would be advantageous to interrupt the Node B's ongoing transmission, allowing access to the channel without further delay. If an interrupt is found to be appropriate a special interrupt message is generated and sent to Node B. The gap in between framelets allows more than enough time for this.

On receipt of an interrupt message a Node B generates and transmits an *interrupt acknowledgment* and ceases its own transmission. The interrupt acknowledgment lets Node A know the interrupt was successfully received by Node B and that it has conceded the channel. In effect the interrupt and acknowledgment messages act as an RTS-CTS handshake and prevent collisions in the case of multiple nodes attempting interrupts simultaneously. This interrupt mechanism can be seen in Fig. 3.

3) *High Priority Interrupt*: Most WSNs have to support several different types of message. As well as needing to exchange control information there are readings collected from a variety of sensors and messages with time constraints attached. All of these messages may be of different significance and it would be useful to prioritize their transport within the sensor network accordingly. A scheduling mechanism could be employed to order the transmission of messages in order of priority, but this would require knowledge of all messages even before they are generated. This is not always possible, particularly for event driven systems, and would most likely require significant overhead in terms of control message.

In addition, even if a high priority message is scheduled immediately, the node might find a (lower priority) message is currently being sent by a neighboring node. This results in the node *backing off* until the channel becomes available and can lead to long delays. These delays can be compounded in multi-hop networks and can become unacceptably large. The interrupt mechanism described in the previous section can be used to gain access to the channel promptly if a high priority message needs to be delivered.

An example of this scheme can be seen in Fig. 3. Node B has a high priority message to send, however before




Fig. 3. High Priority Interrupts.




Fig. 4. Opportunistic Aggregation

transmission it must first check the channel is clear so it performs a presend listen. It finds the channel is currently occupied by Node A with a lower priority message. Since Node B's delivery is more urgent it generates a priority interrupt and sends it to Node A during the gap in between framelets. Upon receipt of the interrupt Node A replies with an interrupt acknowledgment and backs off and reschedules its transmission for a later time. Once Node B receives the interrupt acknowledgment it knows the channel is now clear for it to begin sending.

Node B could in turn be interrupted by another node, say Node C, if Node C's transmission was found to more important. In effect it is a method of arbitration and ensures the node with the highest priority message always gains access to the channel. This Priority interrupt scheme was first described in [22].

4) *Opportunistic Aggregation*: The interrupt concept described can also be used to introduce an aggregation scheme to the framelet protocol.

A node (called Node A for example) with a message to send must first sample the channel for a fixed duration to ensure that the channel is clear. During this presend listen phase any messages on the channel will be intercepted. Should such a message be intercepted from a neighbouring node, Node B say, Node A can compare message headers to determine if they have a common destination. If it is found that the intended message recipients are not the same Node A simply backs off and attempts transmission at a later stage. On the other hand if both messages have the same destination and there is sufficient space to fit their respective payloads into a single message Node A generates and sends an interrupt to Node B. Upon receipt of the interrupt Node B replies with an interrupt acknowledgment and ceases transmission, handing over the channel and its payload to Node A. This is known as packet stuffing and is one of the simplest forms of aggregation. Other forms of aggregation can also be used with this scheme and are described in [7].

Aggregation can increase data throughput without modifying the network's duty cycle and is particularly useful in systems that may experience bursty traffic patterns.

5) *Adaptive Duty Cycle*: While the Opportunistic Aggregation mechanism described in the previous section can be very effective under certain conditions, it is not always possible to aggregate messages. To take this into account FrameComm includes a traffic aware Adaptive Duty Cycle. For example should a node detect an increase in the number of messages received over a certain period it could reduce its sleep interval to accommodate the extra traffic. A subsequent drop in traffic would also be detected allowing the node return to its default duty cycle. These features are highly configurable and can be tailored to the specifications of individual applications.

In addition to traffic awareness certain message types can be used to trigger a change in duty cycle. For the BAN application an alert message sent when a position change is detected indicates a high priority network state where deliver delays must be minimised. To comply with this requirement the duty cycle of the relay node is modified.

V. EVALUATION

Both the Tmote Sky and DSYS25z platforms are equipped with ChipCon's CC2420 radio transceiver, however they employ different microcontrollers. Since none of the code is microcontroller specific and some of the hardware components of the BAN are still under development it was decided to evaluate the proposed software on a small laboratory network comprised of Moteiv Tmote Sky nodes running TinyOS 2.0.2. Each node was programmed to generate data consistent with that produced by the corresponding sensor.

A. Topology

The network consists of two nodes equipped with accelerometers (called Leg and Torso), one node attached to a continuous blood pressure monitor (called Beat-to-Beat), one node connected to an ECG (called ECG) as well as a central forwarding node and a sink or basestation. The sink is comprised of an *always on* node connected to a PC and also acts as a gateway to external networks such as the Internet or a GSM phone network. The nodes are duty cycled to conserve energy and periodically samples data and does some preliminary analysis. The collected data is then sent to the central node from which it is forwarded to the sink, where it is timestamped and logged for further analysis.

The node equipped with the pulse oximeter was excluded because it generates data at a very low rate (one message every 15 minutes) making its impact on network behaviour negligible.

The network has two basic operating states, denoted by GREEN and RED. The default state for the system is GREEN and indicates that the patient has not recently had any major changes in position. The nodes operate on a low duty cycle and those equipped with sensors have relatively low sample rates.

As soon as a position change is sensed by either of the accelerometers, eg sitting to standing, the system transitions to a RED state for a period of two minutes (the time during which a sudden drop in blood pressure is most likely). The RED state

indicates a higher possibility of a sudden blood pressure drop and requires an increased sample rate for both the Beat-to-Beat and ECG nodes. During this time it is vital that any emergency is detected by the sink as quickly as possible so an alert can be raised without delay. To meet this requirement the central node modifies its duty cycle to minimise the sleep interval. Only Beat-to-Beat and ECG readings are of interest in this state, so all other nodes suspend sampling for its duration.

B. Experiment

The purpose of the experiment is to measure the delivery times, t , and the loss rate, r , for each state under typical conditions. BAN nodes generate data consistent with the corresponding sensor and send it to the central node from where it is relayed to the sink. In order to conserve energy and reduce traffic and contention, a number of samples are aggregated into a single message.

Each node begins the experiment in a GREEN state, with a sample rate and duty cycle as shown in Table I. A button on each accelerometer node emulates a change in the patient's stance. When a change is detected the node broadcasts a command message indicating the alert. The payload of the broadcast message contains the duration of the alert, the alert initiator and the sensor readings that triggered it. Once an alert message is received the nodes change sample rate and duty cycle according to the RED state settings in Table I for the amount of time specified in the message. For this evaluation each run lasts 60 minutes, during which time five alerts are triggered, with each alert lasting two minutes. The experiment was run three times.

The purpose of the experiment is to determine whether or not the features of FrameComm, namely, message priorities, opportunistic aggregation and an adaptive duty cycle, make it suitable for use in applications such as falls assessment.

C. Metrics

1) *Sensor Readings*: The purpose of the wireless BAN is to collect data from the various sensors to detect and record changes in the patient's physical condition over a given period of time. As such the data from each sensor needs to be sent to the base station where it is recorded along with a timestamp. Once logged this information is available for analysis by health care professionals. While these sensor readings are of interest to the intended users of the system, they hold little value in terms of network performance and operation. For the purpose of this evaluation only the output from the accelerometers affects the context of the system and how it operates.

2) *Delivery delay*: Introducing duty cycled communication to the BAN leads to an increase in delivery delay. Since the system is dealing with high priority situations where this delivery delay needs to be minimised it is important for it to be recorded for each stage.

The main measurement of interest is the node-to-node delivery latency, t , of messages. However, without a common time source or complex synchronisation mechanism it is difficult to measure this latency. With this in mind it was decided to




Fig. 5. Example Timeline for Test Case 1

measure the latency locally on each node by comparing the time the node decides to schedule a message for transmission, t_1 , and the time the node receives an acknowledgment for that message, t_2 , ($t = t_2 - t_1$). More specifically, within the TinyOS implementation the period between the application calling a *send* and receiving a *sendDone* signal is measured. There is some overhead associated with a message being sent down the TinyOS communication stack and the *sendDone* being delivered, however this time does not vary between experiments or across nodes. The message size is constant throughout and the nodes are not carrying out any other operations or activities that could affect these times. As a result the only variable is the time taken to access the channel and deliver the message. The number of packet losses, r , (given in percent) is another area of interest and is also ascertained.

D. Test Case 1 - Beat-to-Beat

The BAN in this scenario is as described previously but does not include the ECG node. All the relevant configuration parameters can be seen in Table I.

In the RED state the readings from the Beat-to-Beat node are of the highest priority, its sample rate is increased to give more detail and allow faster detection of blood pressure changes. The central node also modifies its duty cycle to minimise sleep time so messages can be delivered with as little delay as possible. While in the RED state the Leg and Torso nodes suspend sampling.

The timeline of one of the experiment runs can be seen in Fig. 5. In this case the alerts were triggered after 5, 27, 35, 43 and 49 minutes. The change in state can clearly be seen by the increase in messages per minute received from the Beat-to-Beat node along with the reduction of those received from the Leg and Torso nodes.

Fig. 6 shows the latency, t , of each node in both states. The expected latency in such a duty-cycled network is given by $\frac{P}{2}$, where P denotes the duty cycle period. With FrameComm a 2% duty cycle equates to $P = 600ms$, giving an expected average latency of just over $300ms$ (includes sending overheads). The average latency achieved for all nodes while in the GREEN state is $280ms$, in the case of the Torso node it is significantly less. This can be explained by FrameComm's

Node No.	Sensor/Role	Sample Data Size	Samples per Message	Green State		Red State	
				Duty Cycle	Sample Rate	Duty Cycle	Sample Rate
1	Central	n/a	n/a	2%	n/a	90%	n/a
2	Accelerometer - Leg	16 bits	5	2%	1000 ms	2%	n/a
3	Accelerometer - Torso	16 bits	5	2%	1000 ms	2%	n/a
4	Beat-to-Beat	16 bits	5	2%	2000 ms	2%	200 ms
5	ECG	8 bits	10	2%	1000 ms	2%	15 ms

TABLE I
NODE AND SENSOR CONFIGURATION FOR EXPERIMENT

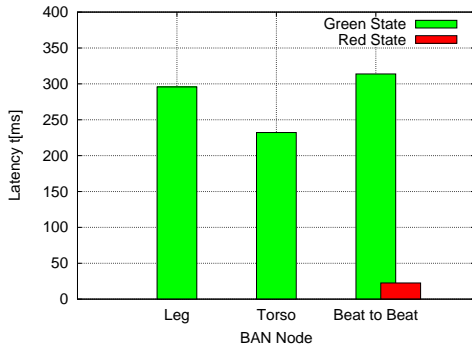


Fig. 6. Latency, r , for Test Case 1

Node	Latency, t		Packet Losses, r
	Average	Std Dev	
Beat-to-Beat (RED)	22.59	4.41	-
Beat-to-Beat (GREEN)	313.76	174.07	0.07
Torso (GREEN)	232.18	181.97	0.11
Leg (GREEN)	295.70	182.03	0.38

TABLE II

TEST CASE 1 RESULTS (ALL PACKET LOSSES FOR BEAT-TO-BEAT NODE ARE INCLUDED IN GREEN VALUE)

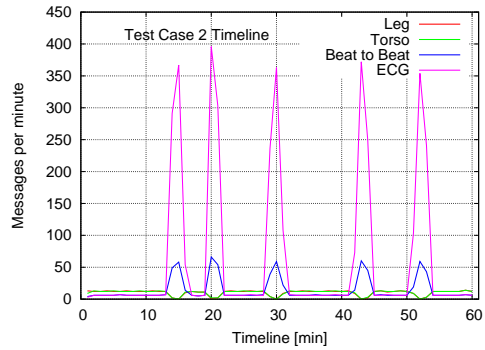


Fig. 7. Example Timeline for Test Case 2




Fig. 8. Latency, r , for Test Case 2

opportunistic aggregation feature, here the Torso node had more of its messages aggregated into those of other nodes.

The effect of the adaptive duty cycle is evident in the difference in latency for the Beat-to-Beat node in each state. The average latency when in the RED state is just $23ms$ compared to $314ms$ for GREEN. This improvement is a result of the change in duty cycle of the central node. By reducing the sleep interval messages can be delivered to the sink more quickly, it also reduces latency and increases bandwidth available.

Table II gives details of the results for each node. The table shows that significant deviation exists in the latency for each node, this is due to the duty cycle of the central node and gives delivery times of anywhere between $0ms$ and $588ms$ ($P = \Delta_0 + \Delta$, given $\Delta = 12ms$ and $P = 600ms$, we get $\Delta_0 = 588ms$), excluding sending overheads. Packet losses, r , are very low and indicate that FrameComm is reliable and suitable for applications such as this.

E. Test Case 2 - Beat-to-Beat and ECG

This scenario is similar to that of Section V-D, but with the addition of the ECG node. Again the configuration is given in Table I, with three 60 minute runs, each of which has five RED alerts triggered.

A timeline for one of the runs is represented in Fig. 7. The five emulated position changes in minutes 13, 19, 28, 41 and 50 can be clearly seen by the spikes in message rate for both the Beat-to-Beat and ECG nodes. Once the event has been detected the Leg and Torso nodes again suspend sampling and the sample rate and those of the Beat-to-Beat and ECG nodes are increased. The higher message generation rate of the ECG compared to that of the Beat-to-Beat is also evident.

Fig. 8 shows the average latency for each node in both GREEN and RED. For all nodes the average latency is close to the expected $\frac{P}{2}$, where again $P = 600ms$. Similar to the previous scenario the Torso node benefits from FrameComm's

Node	Latency, t		Packet Losses, r
	Average	Std Dev	Average
ECG (RED)	40.85	93.46	-
Beat-to-Beat (RED)	85.02	165.64	-
ECG (GREEN)	298.27	186.77	0.84
Beat-to-Beat (GREEN)	342.52	224.70	3.07
Torso (GREEN)	246.66	187.55	0.05
Leg (GREEN)	297.34	174.63	0.27

TABLE III

TEST CASE 2 RESULTS (ALL PACKET LOSSES FOR BEAT-TO-BEAT AND ECG NODES ARE INCLUDED IN GREEN VALUE)

opportunistic aggregation feature. The increased traffic from the ECG node leads to higher latency for the Beat-to-Beat node than that of Fig. 6, however it is still well below the GREEN state value. More detail can be seen in Table III. This shows a slight increase in Packet Losses, this again is due to the increased traffic and associated contention from the ECG node and indicates the network is reaching the limits in terms of available bandwidth, even with the modified duty cycle of the relay node.

VI. CONCLUSION & FUTURE WORK

Wireless BAN technology is emerging as a significant element of next generation healthcare services. In this paper a wireless BAN is presented within the context of falls assessment among elder patients (65+ years). For a BAN to operate effectively it needs to be context aware and react accordingly. The FrameComm protocol, employed by this BAN, is designed to ensure that patient vital signs (i.e. data packets) are communicated under a variety of network loads and conditions.

The BAN is evaluated and the results presented clearly demonstrate the effectiveness of FrameComm with low duty cycles unless an event of interest (i.e. position change) occurred. When an event of interest did occur FrameComm was capable of effectively managing and coordinating data packets with little or no loss of data. Such an outcome with schedule based protocols alone would not be possible.

It is planned to perform a comprehensive evaluation of the BAN in the near future, including repeating the experiments described in this paper. This evaluation will involve wearing the BAN for extended periods of time in various environments, during which time sensor readings will be collected and logged for analysis. Multipath fading, transmitting messages through the human body and interference from Wi-Fi networks and households appliances, such as microwaves, are just some of the factors that can affect reliability. As a result any change in the number of packet losses is of particular interest.

Not all medical episodes can be clearly identified through patient movement. The next version of the FrameComm will build on the current data management techniques to include a richer set of context triggers.

Acknowledgments: Mr. O'Donovan is supported by Microsoft Research through its European PhD Scholarship Programme and the EMBARK Initiative of the Irish Research

Council for Science, Engineering and Technology. Part of the research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 224282. This work is partially funded by the Tyndall National Institute through the Science Foundation Ireland's National Access Programme.

REFERENCES

- [1] J. O'Donoghue and J. Herbert. Profile based sensor data acquisition in a ubiquitous medical environment. In Proc of 4th Annual IEEE Pervasive Computing and Communications Workshops (PerCom), 2006.
- [2] Frost and Sullivan. European Remote Patient Monitoring Markets, June 2008
- [3] Frost and Sullivan. Strategic Opportunity Assessment for Wearable Wireless Patient Monitoring Markets in Europe, October 2005
- [4] K.A. O'Connor, M. O'Connor and E. Moriarty. Cork North Lee and Cork South Lee Local Health Areas: Local Falls Services Mapping Exercise & Strategy Document, 2008
- [5] B. Gannon, E. O'Shea and E. Hudson. The Economic Cost of Falls and Fractures in People aged 65 and over in Ireland. Technical Report to NCAOP/HSE/DOHC, 2007
- [6] SeniorWatch, <http://www.seniorwatch.de/>
- [7] Wearable Technology, Special Issue of the IEEE Engineering in Medicine and Biology Magazine, vol. 22, 2003.
- [8] C. S. Ikehara, E. Biagioni, and M. E. Crosby, Ad-hoc Wireless Body Area Network for Augmented Cognition Sensors In Foundations of Augmented Cognition, 2007
- [9] D. Konstantas, A.T. van Halteren, R.G.A. Bults, K.E. Wac, V.M. Jones and I.A. Widya. Body Area Networks for Ambulant Patient Monitoring Over Next Generation Public Wireless Networks. In: 14th IST Mobile and Wireless Communications Summit, 2004
- [10] J.Y. Khan, M.R. Yuce and F. Karami. Performance Evaluation of a Wireless Body Area Sensor Network for Remote Patient Monitoring. Proc 30th IEEE International Conference on Engineering in Medicine and Biology Society (EMBS), 2008
- [11] M. Sukor, S. Ariffin, N. Fisal, S.K.S. Yusof, and A. Abdallah. Performance Study of wireless Body Area Network in a Medical Environment. Proc. 2nd Asia International Conference on Modeling & Simulation (AICMS), 2008
- [12] J. O'Donoghue and J. Herbert. Data Management System: A Context Aware Architecture For Pervasive Patient Monitoring. Proc. 3rd International Conference on Smart Homes and Health Telematic (ICOST), 2005
- [13] J. Benson, T. O'Donovan, U. Roedig, and C. Sreenan. Opportunistic Aggregation over Duty Cycled Communications in Wireless Sensor Networks. Proc. IPSN Track on Sensor Platform, Tools and Design Methods for Networked Embedded Systems (IPSN/SPOTS), 2008.
- [14] Ahmed Faheem. Wireless Body Area Sensor Network
- [15] W. Ye, J. Heidemann and D. Estrin. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. Proc. 21st IEEE Conference of Computer and Communications (INFOCOM), 2002
- [16] A. El-Hoiydi, J.-D. Decotignie, C. Enz and E. Le Roux. WiseMAC, an Ultra Low Power MAC Protocol for the WiseNET Wireless Sensor Network. Proc. 4th ACM Conference on Embedded Networked Sensor Systems (SenSys), 2006
- [17] W. Ye, F. Silva and J. Heidemann. Ultra-Low Duty Cycle MAC with Scheduled Channel Polling. Proc. 4th ACM Conference on Embedded Networked Sensor Systems (SenSys), 2006
- [18] Analog Devices ADXL202, <http://www.analog.com/>
- [19] Finapres Medical Systems Portapres, <http://www.finapres.com/>
- [20] B. O'Flynn, P. Angove, J. Barton, A. Gonzalez, J. O'Donoghue and J. Herbert. Wireless Biomonitor for Ambient Assisted Living. Oral Presentation at Conference on Signals & Electronic Systems (ICSES), 2006.
- [21] A. Barroso, U. Roedig, and C. J. Sreenan. Use of Framelets for Efficient Transmitter-Receiver Rendezvous in Wireless Sensor Networks. Proc. 5th IEEE Workshop on Wireless Local Networks, 2005.
- [22] T. O'Donovan, J. Benson, U. Roedig and C.J. Sreenan. Priority Interrupts of Duty Cycled Communications in Wireless Sensor Networks. Proc. 3rd IEEE Workshop on Practical Issues in Building Sensor Network Applications (SENSEAPP), 2008