

Remote Physiological Monitoring of First Responders with Intermittent Network Connectivity

Jingyuan Li^{*} Tejaswi Tamminedi[†] Guy Yosiphon[†] Anurag Ganguli[†]
Lei Zhang[†] John A. Stankovic^{*} Jacob Yadegar[†]

ABSTRACT

First responders have been observed to be at an increased risk of cardio-vascular diseases compared to the general population with a high percentage of cardiac events occurring during mission execution. Continuous physiological monitoring during missions can be effective in reducing the number of fatalities. Real-time physiological data such as ECG can be collected using sensors worn on the body. This sensor data can be processed on the body itself or can be communicated over an ad hoc wireless network to the incident command center or base station located near by. First responder missions often take place inside building structures where network connectivity is intermittent. Intermittent connectivity can lead to loss of critical physiological data or delay in that information reaching the base station. Hence, some amount of local processing is needed in order to limit the data that is communicated. In this paper, we introduce a novel Hidden Markov Model based myocardial infarction detection approach. The fidelity of this approach can be adapted based on the processing power available. We present a peer-to-peer networking protocol for communication over disrupted networks. A low fidelity classifier is used to perform local processing and assign priorities to the data based on its criticality. It is complemented by a disruption-aware epidemic forwarding protocol for transferring first responder's physiological data to the base station. With prioritized epidemic forwarding and buffer eviction policy, our protocol increases packet delivery ratio and reduces networking delay when end-to-end route disruption occurs. Finally, we report the effect of network disruption on myocardial infarction detection rate and latency of detection and the improvements achieved by our protocol.

^{*}Department of Computer Science, University of Virginia, Charlottesville, VA, {jl3sz,stankovic}@cs.virginia.edu

[†]UtopiaCompression Corporation, 11150 W. Olympic Blvd, #680, Los Angeles, CA 90064, {tejaswi,guy,anurag,lei,jacob}@utopiacompression.com

1. INTRODUCTION

First responders have been observed to be at an increased risk of cardio-vascular diseases compared to the general population with a high percentage of cardiac events occurring during mission execution [7, 20, 23]. Continuous physiological monitoring during missions can, therefore, be effective in reducing the number of fatalities. One possible system architecture to enable such continuous monitoring is as follows [9]. Physiological sensors, preferably non-invasive, measure real-time physiological data. A portable personal server, on the first responder body, collects the various sensor readings. The personal server performs two additional functions. It performs preliminary sensor data processing and is also responsible for communicating the data to a remotely located incident command center or base station. Additional processing and decision support capability may be available at the base station. In this paper, we address some of the challenges associated with remote physiological monitoring of first responders.

Such a continuous physiological monitoring operation, however, presents several challenges. One challenge that arises is due to the fact that first responder missions often take place inside building structures where network connectivity is intermittent. The intermittent network connectivity may lead to the loss of critical physiological sensor data. A disruption tolerant networking solution is, therefore, necessary to address this problem. However, it is still not guaranteed that all transmissions will reach the base station. Moreover, data may still incur significant latency before it reaches the base station. To alleviate this problem, it is important that the more critical information have a higher chance of successfully reaching the base station and earlier. To identify the more critical data, the personal server must be able to perform preliminary processing on the sensor data and gauge its importance. Cardiovascular conditions can best be detected with Electrocardiography (ECG) and, therefore, the personal server must be able to provide some diagnostic capability using ECG measurements. In this paper, we consider the problem of communicating physiological data from a first responder to a remotely located base station in such a manner so that the latency associated with critical physiological ECG data is minimized and overall probability of detection of an abnormal condition is maximized.

The first set of relevant literature is in the area of automated ECG analysis. Techniques for automated ECG analysis have been studied for more than two decades. Most

automated ECG analysis approaches rely on machine learning techniques [5, 21, 18, 16, 17, 24] and primarily focus on the detection of arrhythmias. ECG analysis using mobile devices [14] has also been studied. Automatic detection of myocardial infarction (MI) using ECG data has not received as much attention as the detection of arrhythmias. To the best of our knowledge, there is very little literature on MI detection using single lead ECG. Documented results using 12 lead ECG [3] obtain results of 53% sensitivity for MI and 83% for Acute MI. Using 12 lead ECG during first responder missions is not practically feasible.

The second set of relevant literature is in the area of Delay or Disruption Tolerant Networking (DTN). Fire fighter missions may take place in challenged environments such as building structures, tunnels, subterranean environments where network connectivity to the base station is intermittent. Without a DTN solution in place, critical physiological data may be lost or incur significant latency. DTNs have been studied in the context of vehicular networks [12], fire fighter networks [22] and wildlife tracking networks [15]. Routing protocols for DTN have been proposed that utilize the knowledge of mobility of the nodes [13] or using an epidemic based approach [8]. In the first responder case, it is difficult to assume that the mobility knowledge is available. Hence, an epidemic based forwarding approach is more suitable for our application.

Our contributions in this paper as follows. First, we propose a novel automated myocardial-infarction detection approach using ECG data. This algorithm can be used in a decision support system located at the incident command center to help provide real time assessment of the first responders' health. Our approach is based on a variant of a Hidden Markov Model (HMM) approach. Our approach uses Duration Hidden Markov Models (DHMMs) which capture the temporal characteristics of the ECG signal very well. To account for the large variability in the ECG waveform among patients suffering from the same cardiac condition, we train multiple DHMMs for each class. The outputs of the multiple DHMMs are combined together to yield one classification result.

To evaluate our approach, training and testing was conducted using data from the Physikalisch-Technische Bundesanstalt (PTB) database on Physionet [1] which contains patient ECG posterior to MI. By using only lead 2 of the ECG we obtained about 70% accuracy in detection with 0% false positives (90% AUC by utilizing parallel DHMMs). A crucial point to note is that patients used for testing were completely distinct from those used in training, akin to real life situations. This is very rarely found in ECG literature due to the issues in addressing high variability between patients and also in severity of the condition. We also provide a lower fidelity version of the above DHMM classifier that can be implemented on the first responder personal server to provide preliminary screening and identify data which is critical.

Second, we present a delay-tolerant networking (DTN) approach to communicate ECG data from the personal server to the base station. Data that is more likely to contain abnormal ECG receive higher priority for communication. In

the case of buffer overflows, lower priority data is evicted first. To compare our approach with the baseline approach without DTN, we conduct simulations using realistic first responder mobility patterns. Prioritized ECG data is generated at each first responder node as the network traffic. We observe that our approach improves the overall performance in terms of higher average packet delivery ratio, higher packet delivery ratio for the abnormal ECG data and lower latency. We also evaluate the effect of buffer size on the performance metrics.

The paper is organized as follows. Section 2 contains a description of the first responder physiological monitoring system architecture. Section 3 describes the HMM based approach for MI detection. Section 4 describes our disruption-aware forwarding protocol. Section 5 describes the experimental set up. Section 6 contains the experimental results and discussions. Finally, Section 7 contains the conclusions and pathways for future work.

2. SYSTEM ARCHITECTURE

In this section, we describe the overall architecture of a prototypical first responder physiological monitoring system. Each first responder wears a number of physiological sensors on his or her body. A vital sensor to monitor abnormal cardiac conditions is the ECG sensor. These sensors communicate to a personal server also located on the body of the first responder. The personal server can be a mobile device with some processing and communication capabilities. The personal server can be used to communicate over a wireless interface to other first responders in a peer-to-peer manner or to the incident command center (base station). Please refer to Figure 1 for a graphical illustration of the system architecture.

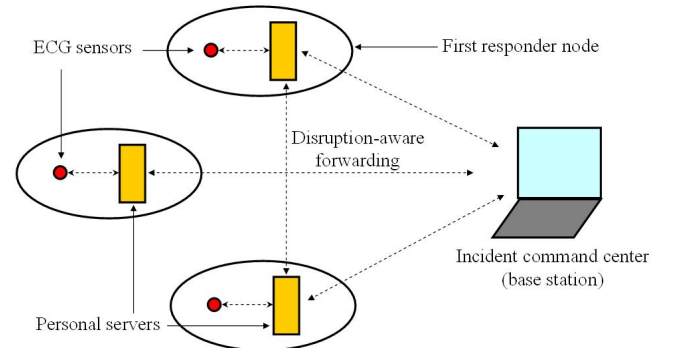


Figure 1: System architecture of a first responder remote physiological monitoring system

The personal server collects data from the ECG sensor, performs preliminary diagnosis to gauge the importance of the data, and communicates the data along with a measure of its importance using a disruption-aware epidemic forwarding protocol. The incident command center or the base station has a *Decision Support System (DSS)* to help further analyze the ECG data and assess the health status of a first responder. In principle, a copy of the DSS can also reside on the personal server. However, high fidelity algorithms cannot be run on the personal server because of its limited processing

power. The limited processing power will inject increased latency which can be disastrous in the event that a first responder suffers a heart attack. In this paper, we describe a novel high fidelity Hidden Markov Model (HMM) based ECG analysis algorithm for the DSS and a lower fidelity version of the same algorithm for preliminary diagnosis at the personal server. These algorithms are described next in Section 3. The disruption-aware epidemic forwarding protocol is described later in Section 4.

3. HMM BASED MI DETECTION

As described in the previous section, the DSS located at the base station consists of a HMM based MI detection algorithm. A lower fidelity version of the same algorithm is located on the personal server to perform preliminary diagnosis. These algorithms use real time Electrocardiogram (ECG) signals as inputs. The incoming raw ECG signal is filtered and then processed by multiple Hidden Markov Models (HMMs) [19]. The HMM models are trained using real patient data obtained from the Physionet PTB dataset [1]. The data includes ECG data of patients suffering with Myocardial Infarction(MI) and healthy subjects.

This section is organized as follows. Section 3.1 presents a brief overview of MI and its ECG symptoms. Section 3.2 provides an overview of conventional HMM modeling and inference. Section 3.3 describes a generalization of the conventional HMM models that includes an explicit probability distribution for the duration of each state. This extended model with Gaussian distribution of state durations can model accurately the durations of each state in an ECG signal. In Section 3.4, we provide specifics of the implementation of the HMM based classifier followed by a discussion on adjusting the fidelity of the classifier for platform specific use in Section 3.5. Finally, in Section 3.6, we describe priority assignment of ECG data at the personal server before transmission.

3.1 Myocardial Infarction

The primary purpose of our classifier is to identify abnormalities in the ECG that are indicative of Myocardial infarction (MI). Commonly known as a heart-attack, MI is the interruption of blood supply to parts of the heart. The resulting ischemia (restriction in blood supply) and oxygen shortage, if left untreated for a sufficient period of time, can cause damage or death (infarction) of heart muscle tissue (myocardium). Among the diagnostic tests available to detect heart muscle damage are ECG, chest X-ray and various blood tests. Acute, evolving or recent MI causes changes in the ECG wave. A normal ECG wave along with its various segments is depicted in Figure 2. MI causes distortions in some segments such as development of pathological Q-waves and ST segment elevation or depression [2].

3.2 Hidden Markov Models

An Hidden Markov Model (HMM) is a probabilistic model of a temporal process where the assumption is that a sequence of observations is generated by an unobserved or hidden Markovian state variable. HMMs have been widely used for modeling and classification of time series data such as speech [11] and ECG [19].




Figure 2: ECG waveform and its component segments

Consider a stochastic process which may be described at any time as being in one of N distinct states, S_1, S_2, \dots, S_N . At regularly spaced discrete times, the system undergoes a change of state according to transition probabilities associated with the current state (hence a Markovian process). These state transition probabilities can be represented as follows:

$$a_{ij} = P(q_t = S_j | q_{t-1} = S_i), \quad 1 \leq i, j \leq N \quad (1)$$

where q_t is the state at time t , and $P(q_t = S_j | q_{t-1} = S_i)$ is the probability that the state at time t is S_j given that the state at time $t-1$ is S_i . A probability distribution of the observation is associated with each state S_i denoted by B_i . The set of B_i 's is denoted by B . In case the observations are continuous (as in ECG data), the observation probability distributions can be any continuous distribution such as Gaussian. The probability of the initial state being S_i is distributed according to the density π . The entire HMM can, therefore, be conveniently represented as follows:

$$\Lambda := (A, B, \pi). \quad (2)$$

Training an HMM model can be achieved by maximizing the conditional probability $P(\Lambda|O)$ where O is a set of observation sequences. The Baum-Welch algorithm [4], which is a special case of the expectation-maximization algorithm [6], is a well known procedure for training HMMs. Once a set of possible models have been trained, say $\{\Lambda_1, \Lambda_2, \dots\}$, the classification problem is to identify the most likely model for a given set of observations. This is done by maximizing $P(O|\Lambda_k)$ over all k . The procedure for evaluating $P(O|\Lambda)$ is the Forward algorithm [19], which has linear complexity in the number of observations and states.

3.3 Explicit-duration HMMs

A major weakness of conventional HMMs is the modeling of state duration [19]. The inherent duration probability of a conventional HMM to remain in a given state is an exponential distribution. This exponential state duration density is inappropriate for most physical signals [5]. Instead, a better choice is an HMM with explicit duration density in some analytic form. In this type of HMM model the self-transition coefficients a_{ii} are set to zero and an explicit duration density is specified. A sequence of states can be sampled from an explicit-duration HMM according to the following sequence of operations.

1. Chose initial state q according to the initial state distribution π .
2. Choose the duration d of the state q according to the duration probability density $p_q(d)$.
3. Sample a set of d independent observations O_1, O_2, \dots, O_d according to the observation density b_q .
4. Choose the next state according to the transition probabilities a_{ij} where $i \neq j$.

The training and inference algorithms of regular HMMs are applicable for explicit-duration HMMs. However, the computational cost of these models is much higher than regular HMMs. The time complexity is increased by a factor of D^2 where D is the maximum duration within each state.

3.4 HMM classifier implementation

For the MI detection application, the hidden states chosen are the identifiable parts or sub-parts of an ECG waveform, namely ISO, P,PQ, Q, R1, R2, R3, S1, S2, S3, ST and T (figure 2). A Gaussian observation distribution model $B_j = \{\mu_j, \Sigma_j\}$ (the mean vector and covariance matrix) is used to model the continuous ECG signal. Each hidden state is associated with a different mean and variance and therefore models a different segment of the ECG wave. We choose a Gaussian explicit duration density since a reasonable bound on the maximum duration can then be obtained which is a function of the variance of the Gaussian density.

The Baum-Welch algorithm can train any unconstrained HMM model given the set of observations. However, the Baum-Welch algorithm is notorious for getting trapped in local minima. Restarting from random initial values can help only in limited way since the combinations of different states increase exponentially in the number of hidden states. Moreover, if the transitions are not constrained then the number of different transitions matrices is exponential in the square number of states. Clearly, there is a need to constrain the number of free parameters in order to achieve acceptable classification performance in reasonable time.

Therefore, for our implementation, we constrain our explicit-duration HMM model to allow only transitions between subsequent states, i.e. from i to $i + 1$ or from N to 0. Recall that the hidden states represent known segments of a single ECG wave such as the P, Q, R, S, T waves . The last state (state T) returns to the first state (ISO) in order to begin a new heart wave. Figure 3 depicts the HMM transitions of our ECG model.




Figure 3: HMM states and transitions: The HMM states correspond to the segments in ECG wave (P, Q, R, S, T waves and the intervals in between

To train the HMM classifier we use data from the online

PTB data set [1]. Training requires providing initial conditions for the different model parameters. As shown in Figure 3, the transition probability matrix is known from the structure of the HMM. The unknown model parameters to be learned through training include the mean and variance of the Gaussian duration probabilities and the mean and variance of the Gaussian observation probabilities. Instead of providing random initial values for the model parameters, we use values from a single ECG heart wave. The heart wave segments are manually extracted. For the duration probability Gaussian, we choose the duration of a segment as the mean with variance equal to half of the mean. The mean and variance of the ECG samples that belong to the segment are used as the mean and variance of the observation probability Gaussian. This training approach was observed to significantly improve the overall training performance, as in most cases the subsequent training on the unlabeled ECG did not modify the model's coefficients by considerable amount.

Now, given a test ECG wave, the HMM classifier provides a segmentation of the signal into the different ECG components. Please refer to Figure 4 for an example of the segmentation. This is a major advantage of HMM modeling over other classifiers such as Support Vector Machines (SVM) or artificial Neural Networks (ANN) that cannot provide insightful interpretations of the input.

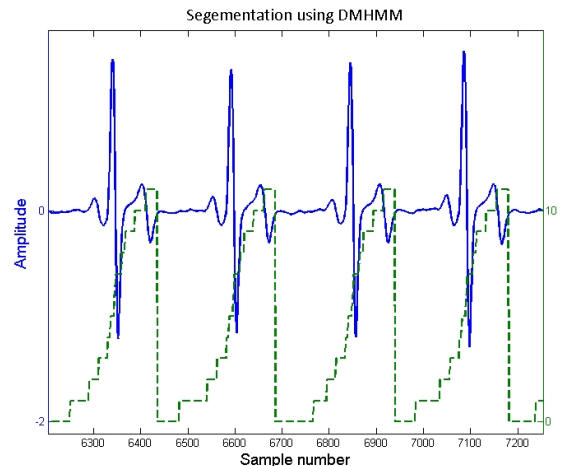


Figure 4: Segmentation of ECG signal according to the HMM states. The dashed line depicts the inferred HMM states - 0 corresponds to ISO and 11 corresponds to T.

We also observed that there is a lot of variability in the ECG from one patient to another even though they might be suffering from the same heart condition. To account for this, we train an HMM model for each subject in the PTB data base. At the end of this training, we have as many trained HMM classifiers as there are patients in the data base. The classification decision of an unlabeled ECG signal is made according to the HMM model that provides the highest likelihood. This scheme is similar to nearest neighbor algorithms where the distance metric is the likelihood function. If we incorporate as many HMM models as pos-

sible the accuracy will increase since the space coverage is increased. However, each additional HMM model increases the computational complexity and thus there is a limit to the number of HMM models that can be evaluated in parallel. Figure 5 shows the accuracy results for different number of HMM models.

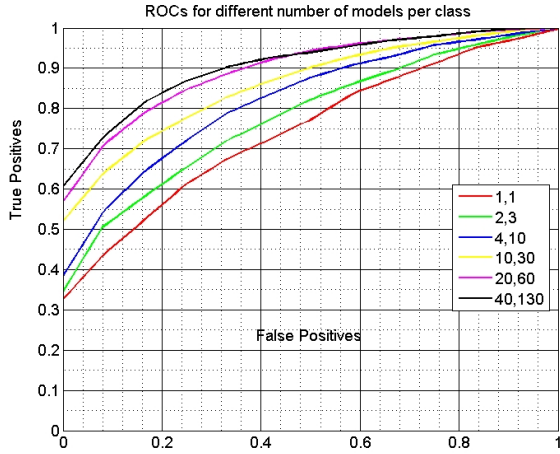


Figure 5: Receiver Operating Characteristic (ROC curve) of the HMM classification over the PTB dataset. Each line represents the results of a classifier comprised of a different number of HMM models. For example, the yellow line represents the ROC of a classifier that is comprised of 10 healthy subject models (picked randomly out of 50), and 30 MI subject models (out of 148).

The factor of D^2 in the computational cost of explicit-duration HMM inference algorithm, is the main reason why small low-powered devices (such as PDAs) cannot process multiple explicit-duration HMM models. Hence, to perform preliminary diagnosis on personal server, we need a classifier that uses fewer number of explicit-duration HMM models. The DSS at the base station on the other hand can use a greater number of models.

3.5 Adapting classifier fidelity based on processing power available

As shown in Figure 5, the HMM classifier performance improves with increasing number of models that it is comprised of. This improvement in performance comes at the cost of increased computational burden. Since the DSS located at the incident command center or base station is expected to have adequate computational resources, a high fidelity classifier should be used there. Such a classifier, if run on the personal server with limited resources, will add a large latency to the overall decision making. Still some amount of preliminary screening is needed at the personal server to assign priority or an importance level to the data. This is necessary because in the event of network disruptions, critical data should be given priority during transmission and also when deciding which data to keep in the event of a buffer overflow.

As shown in Figure 6, we build two classifiers of different fidelities. The lower fidelity classifier is for use on the personal server. The higher fidelity classifier is for use in the DSS at the base station. The difference in the two classifiers is the number of models that they are composed of. The lower fidelity classifier is based on just two models - normal and abnormal. The higher fidelity classifier consists of multiple normal and abnormal models.

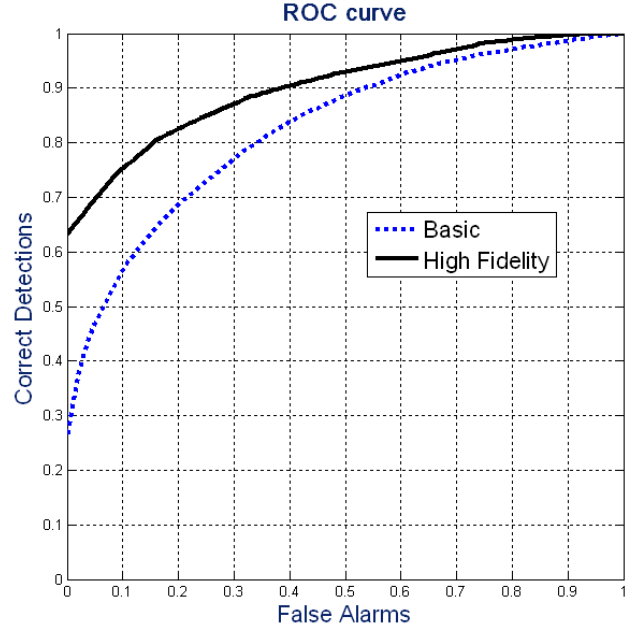


Figure 6: Comparison of ROCs for Basic and high fidelity classifiers

3.6 Priority assignment

The classifier on the personal server assigns a priority to the data being broadcast by it. This priority reflects the chance that the data might contain abnormal physiological measurements that need to be re-interpreted by the base station DSS. The higher the probability of abnormality the higher the priority assigned to it. For the first responder scenarios, the data being transmitted are segments of ECG data and the abnormality to be detected is the occurrence of a Myocardial Infarction (MI). So, higher the chance of the ECG containing an indication of MI higher the priority assigned to it. As described in the previous section, the classifier at the personal server consists of two models, one corresponding to normal and the other to abnormal heart condition. Given a ECG wave, the classifier outputs two likelihoods, one is the likelihood of the ECG segment being normal and the other is the likelihood of the ECG being indicative of MI. The ratio of these two likelihoods, likelihood of MI over normal, called the likelihood ratio (LR) depicts a rising chance of MI with its increase. The priority assigned to the ECG segment needs to be an increasing function of the likelihood ratio and for our simulation purposes we take the log of the likelihood ratio as the priority; see Figure 7.

4. DISRUPTION-AWARE FORWARDING

As described earlier, disrupted network connectivity conditions will lead to data losses, delayed updates of physio-




Figure 7: Priority calculation from ECG

logical parameters to the base station, and may seriously threaten the first responder’s safety. The requirement of prompt and accurate physiological parameter updates under intermittent network connectivity conditions motivates us to employ a disruption-aware epidemic forwarding protocol to help first responders send critical physiological event updates to the base station as quickly as possible. Our forwarding protocol is composed of three basic building blocks: disruption detection, prioritized epidemic forwarding and eviction policy.

4.1 Disruption Detection

First responders need to adjust their forwarding strategy given different network connectivity conditions (connected, or disrupted end-to-end forwarding path) for best packet delivery. Disruption of a forwarding path is caused by first responder mobility and limited communication range of their radio devices. Disruption of end-to-end connectivity is a knowledge that can be learned from failure/null return of a route discovery procedure. In this paper, we do not go into the details of disruption detection solutions; rather we assume that all first responders’ radio devices have the capability of disruption detection, i.e., these devices are disruption aware.

4.2 Prioritized Epidemic Forwarding

When first responder’s radio devices are connected to the base station via a multi-hop forwarding path, physiological data can be uploaded with conventional multi-hop forwarding protocols. However, first responder mobility will lead to end-to-end route disruptions. When a disruption occurs, normally first responders will have to drop the data, or buffer all the data until an end-to-end route is re-established at some later point in time.

However, the conventional drop or buffer-and-wait solution will inevitably cause the loss of important physiological data or delay the update of critical physiological events. To alleviate these problems, we propose to use prioritized epidemic forwarding when first responders experience network disruptions.

When physiological data is processed locally at the personal server, each block of data is assigned a priority. A block of data is usually continuous in time and associated with certain physiological events. The higher the likelihood that the data is associated with an adverse cardiac condition such as MI, the higher the priority it gets. Data blocks marked with priorities are buffered locally during network disruption. Meanwhile, the first responder’s radio device searches for one-hop neighbors, selects a block of data with highest priority, and broadcasts the data to its one-hop neighbors.

Neighbors overhearing the high priority data try to put the data in their local DTN (Disruption-Tolerant Networking) buffer, and upload the data later when an end-to-end route is re-established. This process is called spray-and-wait; this approach, high priority data is replicated and carried by multiple forwarders, and thus the probability of delivery is increased.

4.3 Eviction Policy

As we discussed in the previous section, each first responder device maintains two data buffers: (i) a local data buffer is used to store data that has already been processed by the local HMM classifier but not yet sent to the base station, and (ii) a DTN buffer is used to store high priority data that is overhead from disconnected one-hop neighbors. Since each first responder device only has a finite buffer size, when buffers are completely filled up, certain data has to be evicted to make room for more important ones.

Since our application is remote physiological monitoring, what the users care most is the critical physiological events, such as possible heart attacks. Thus our eviction policy is also priority based, to make sure high priority data will always get delivered. When new data arrives, the first responder node will search the buffer to find free buffer space to accommodate the data. If no free space is found, a block space with lowest priority data is returned. Base on the relative value of the priorities, the first responder device will decide whether to drop the data that just arrived or evict the low priority data block already in the buffer. In other words, the lower priority data is evicted.

Overall, our disruption-aware forwarding protocol takes into account the relative data priority, to make sure that data associated with critical physiological events are more likely to get delivered when intermittent network connectivity conditions occur. Our protocol also takes advantage of first responder mobility, using epidemic forwarding to replicate the high priority data to several one-hop neighbors, thus increasing the probability of delivery and reducing the delay of delivery for data with high priorities.

5. EXPERIMENTAL SET-UP

In this section, we describe our experimental set-up to evaluate the performance of the disruption-aware epidemic forwarding protocol in conjunction with MI diagnosis using the HMM based approach. We designed simulation experiments to try and faithfully capture the performance of the HMM classifier at the personal server as well as capture the network disruptions as they are expected to occur during first responder missions. This section is organized as follows. In Section 5.1, we describe how we generate the ECG data along with priority levels for use in simulation. In Section 5.2, we describe how we generate representative first responder mobility patterns along with the other specifics of the simulations.

5.1 Prioritized data generation

To perform simulations, we need to simulate the data traffic that is generated at any given node. In our case, the data consists of ECG measurements. ECG sensors can be assumed to generate data at a constant rate. In our case,

however, generating raw data is not sufficient. The data also has to be associated with priority levels based on the preliminary diagnosis of the HMM based classifier at the personal server. To generate such priority assigned data, there are two possible alternatives. The first is to use real ECG data from online databases such as the PTB. However, the limited data would constrain the number and time duration of the simulations that can be run. Moreover, available real data sets contain either healthy patients or patients who have already had a heart attack. In the event that a first responder has a heart attack during mission, the ECG will change from that of a healthy person to that of a person suffering from MI. Such real data is not available. The second alternative is to use synthetic ECG data along with a probabilistic model of the HMM classifier performance. In this paper, we adopt the second alternative. We now provide more details of our approach.

First, we build a probabilistic model of the HMM classifier performance. In other words, we estimate the probability density of the output likelihood ratio (see Figure 7) given that the ECG data corresponds to a healthy subject and the probability density of the output likelihood ratio given that the ECG data corresponds to a patient suffering from MI. Such a model can indeed be built using the real data available from the online PTB database. Please refer to Figure 8 for the probability distributions of the basic classifier to be used at the personal server. These distributions are constructed by computing the log likelihood ratio outputs of the basic classifier for all the available ECG samples of both healthy subjects and MI patients from the PTB data base. Since the log likelihood ratio is defined as the log of the ratio of the likelihood of MI to the likelihood of normal, one would expect that for healthy subjects the log likelihood ratios would tend to be less than zero and for MI patients the log likelihood ratios would tend to be greater than zero. This is indeed the case as can be seen from Figure 8.

Given that a sample ECG wave comes from a healthy patient, one could sample from the distribution shown in Figure 8 (top) to get a value of the basic classifier output. Likewise, given the subject has had MI, one could sample from the distribution shown in Figure 8 (bottom).

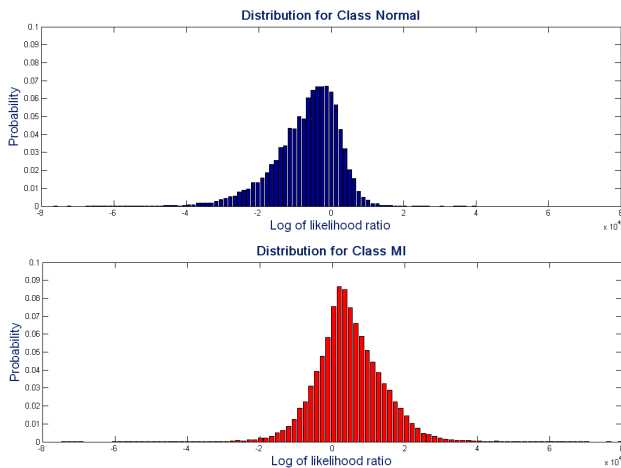


Figure 8: Likelihood ratio distributions

For each first responder node, we generate random synthetic data traffic at the rate of 10 KB per second. This is consistent with that of a prototypical 2 lead ECG sensor operating at 200 Hz. For every contiguous 1 minute interval of data, we also generate the ground truth - a binary number signifying whether the person is healthy or has had MI. Given the ground truth, one can sample from the appropriate distribution in Figure 8 to get the likelihood ratio or priority level for the 1 minute segment. Thus, we can generate surrogate data traffic along with priority level.

5.2 First responder scenario set-up

We generate first responder mobility traces using the CORPS mobility model [10]. CORPS stands for Cooperation, Organization, and Responsive in Public Safety. This model has been developed specifically to capture the unique characteristics of first responder (FR) missions. The model is event-driven - the mobilities are dictated by events such as trapped victim and explosion. The CORPS model is more suited for large scale incidents where tens to hundreds of FRs are involved. Therefore, our simulation results included in the next section are also representative of such incidents. In our simulation, we generate mobility traces for 30 minute first responder missions involving 18 first responders. We assume that one of the first responders suffers a heart attack 15 minutes into the mission. We repeat our experiments for three different scenarios.

For our simulations, we consider a building of size 1000 by 1000 square units and first responder radio communication range equal to 250 units. As shown in Figure 9, there is one base station, node 0, placed at the entrance of the building, node 1 through 18 represent first responder radio devices, the solid lines indicate that two nodes are within radio communication range of each other. First responder locations and connectivity status are updated each simulation step (1 unit time).

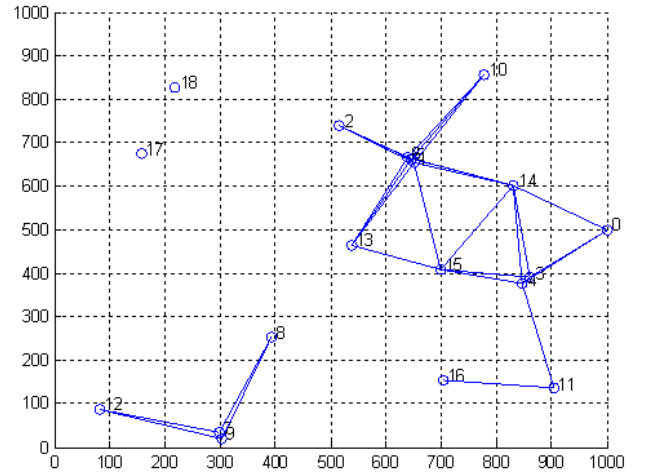


Figure 9: Snapshot of First Responder's Spatial Distributions and Connectivity Status

We assume a constant bit rate (CBR) traffic pattern for our simulation. Data is generated every 60 time units, with

the block size of 60 data units in length and put into the first responder's local data buffer, each block of data is assigned a priority using the method introduced in Section 5.1. Data transmission occurs every time unit, if first responder is connected to base station, 1 unit length of data is moved to base station from the local data buffer (if data buffer is not empty), 1 additional unit length of data is moved from the DTN buffer to the base station (if DTN buffer is not empty); if the first responder is disconnected, 1 unit length of data is broadcast to its one-hop neighbors (if it has any one-hop neighbor), data broadcast from the data buffer will stay in the buffer, but will be marked and will not be broadcast again. Neighbors overhearing the data will store the 1 unit length of data in their DTN buffer. Buffer eviction is based on the policy introduced in Section 4.

6. RESULTS AND DISCUSSIONS

We compare the performance of the baseline operation (without the disruption-aware forwarding feature) with the disruption-aware forwarding protocol with assigned data priority levels. We compare the overall packet delivery ratios, the packet delivery ratios for the abnormal signals (ECG event delivery ratio) and the latency between the heart attack event and the receipt of the abnormal signals at the base station. Connectivity ratio is defined as the fraction of the total simulation time when each firefighter has connections to Base Station. Packet delivery ratio is the total number of non-duplicate packets received from a firefighter over the total number of packets generated during the mission for that individual. Average connectivity ratio and average packet delivery ratio are the average values of these ratios for the 18 firefighters in simulation. Finally, the heart problem detection delay is the time interval between the instant when the heart attack occurs and when an abnormal packet is first received (this includes delays caused by data buffering, data decomposition and reassembly). We also evaluate the effect of buffer size on performance. The results are summarized in Figures 10, 11 and 12. The simulations were run for four different local and DTN buffer sizes shown in the table 1.

Table 1: Buffer assignments for different simulation cases. Each unit of memory represents 1 minute of ECG data.

	1	2	3	4
Local	1	1	5	10
DTN	0	5	10	10

From the figures, we observe that the average packet delivery ratio (APDR) increases when DTN buffer size is increased. Also, the APDR increases when the local buffer size is increased. The improvement plateaus beyond a certain buffer size. This happens because the network connectivity is still sufficient to not cause buffer overflow.

We observe a similar behavior for the ECG event delivery ratio (EEDR) or the abnormal packet delivery ratio. But in this case, the improvements are more pronounced. Recall that abnormal data packets are prioritized for storage and transmission. As local and DTN buffer sizes are increased, a greater fraction of the buffers get occupied by high priority data. Here also we observe a diminishing performance improvement with increasing buffer size which can be explained




Figure 10: Results for scenario 1

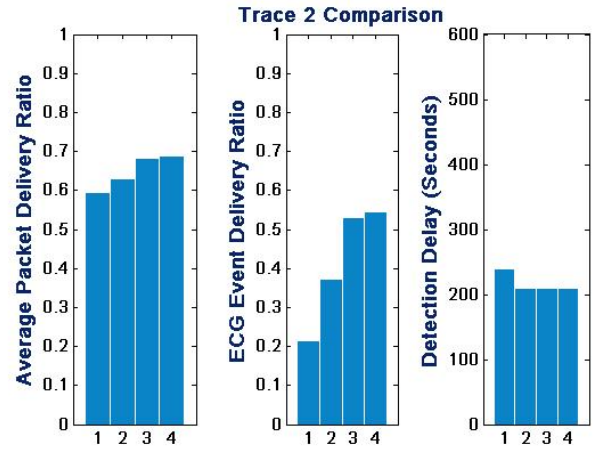


Figure 11: Results for scenario 2

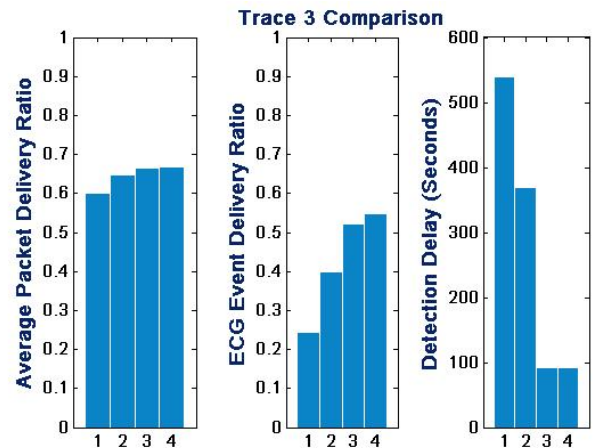


Figure 12: Results for scenario 3

the same way as APDR.

Recall that the heart problem detection delay is the time interval between the instant when the heart attack occurs and when an abnormal packet is first received at the base station. This ideally should only be a function of whether DTN is enabled or not which is shown in two out of the three cases. The reason for this is that an abnormal data can only be evicted by another piece of abnormal data because of the way we assign priorities. However, in practice, the classifier is not perfect and a normal data segment may sometimes be assigned a higher priority than an abnormal data segment. This explains the observation in experiment 3, where the delay decreases when the local buffer size is increased from 1 to 5.

7. CONCLUSIONS

In this paper, we have introduced a novel MI detection approach for use in a first responder physiological monitoring application. We have also introduced a disruption aware epidemic forwarding protocol to communicate physiological data from first responders in the field to the base station or incident command center. We have shown how the MI detector can be integrated with the epidemic forwarding protocol to provide higher packet delivery ratios and lower MI detection times in the face of network disruptions. Intelligent integration of physiological monitoring with networking protocols is important for such wireless health applications where resources such as bandwidth and processing power are at a premium.

The results included in this paper are encouraging but are still only the first steps towards providing a robust disruption tolerant remote physiological monitoring solution for first responders. The protocols and algorithms introduced in this paper need testing in more realistic environments. For example, the ECG used for this study was collected in hospital environments and did not include motion artifacts expected to present in ambulatory ECG monitoring. Also, for the purpose of the network simulations, we assumed a simplified range based model for one-hop connectivity. In real indoor environments, this is seldom the case with many more complex effects such as multi-path fading coming into play. These effects need to be considered for better evaluation of our approach and are part of our future work plan.

8. ACKNOWLEDGMENTS

This effort is funded in part by a Department of Homeland Security (DHS) Phase II Small Business Innovation Research (SBIR) grant. We are grateful to our program manager, Mr. Jalal Mapar, for his support and encouragement. We also thank Ms. Ying Huang of the University of Illinois at Urbana-Champaign for providing us the CORPS mobility model.

9. REFERENCES

- [1] The ptb diagnostic ecg database. Available online at <http://www.physionet.org/physiobank/database/ptbdb/>. Accessed 08/12/2009.
- [2] J. Alpert, K. Thygesen, E. Antman, and J. P. Bassand. Myocardial infarction redefined—a consensus document of the joint european society of

cardiology/american college of cardiology committee for the redefinition of myocardial infarction. *Journal of American College of Cardiology*, 36:959–969, 2000.

- [3] A. Andresen and M. D. Gasperina. An improved automated ecg algorithm for detecting acute and prior myocardial infarction. *Journal of Electrocardiology*, 35(4), 2002.
- [4] L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Ann. Math. Statist.*, 41(1):164–171, 1970.
- [5] D. A. Coast, R. M. Stern, G. G. Cano, and S. A. Briller. An approach to cardiac arrhythmia analysis using hidden markov models. *Transaction on Biomedical Engineering*, 37(9), September 1990.
- [6] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*:1–38, 1977.
- [7] R. F. Fahy. U.s. firefighter fatalities due to sudden cardiac death 1995 - 2004. National Fire Protection Association Fire Analysis and Research Division, June 2004.
- [8] K. Fall. A delay-tolerant network architecture for challenged internets. In *Applications, technologies, architectures, and protocols for computer communications*, 2003.
- [9] A. Ganguli, W. Kaiser, T. Tamminedi, and J. Yadegar. Smart physiological monitoring of first responders. In *SPIE Defense, Security, and Sensing*, Orlando, FL, April 2009.
- [10] Y. Huang, W. He, K. Nahrstedt, and W. C. Lee. CORPS: Event-driven mobility model for first responders in incident scene. In *IEEE Military Communications Conference*, San Diego, CA, 2008.
- [11] N. P. Hughes, L. Tarassenko, and S. J. Roberts. Markov models for automated ECG interval analysis. In S. Thrun, L. Saul, and B. Schoelkopf, editors, *Advances in Neural Information Processing Systems 16*, Cambridge, MA, 2004. MIT Press.
- [12] B. Hull, V. Bychkovskiy, K. Chen, M. Goraczko, E. Shih, Y. Zhang, H. Balakrishnan, and S. Madden. Cartel: A distributed mobile sensor computing system. In *SenSys*, 2006.
- [13] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Applications, Technologies, Architectures, and Protocols for Computer Communication*, 2004.
- [14] Z. Jin, Y. Sun, and A. C. Cheng. Predicting cardiovascular disease from real-time electrocardiographic monitoring: An adaptive machine learning approach on a cell phone. In *31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2009.
- [15] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebrant. In *ASP-LOS*, San Jose, 2002.
- [16] M. Korurek and A. Nizam. A new arrhythmia clustering technique based on ant colony optimization. *Journal of Biomedical Informatics*, 41:874–881, 2008.
- [17] K. Polat and S. Gunes. Detection of ecg arrhythmia

using a differential expert system approach based on principal component analysis and least square support vector machine. *Applied Mathematics and Computation*, 186, March 2007.

- [18] G. K. Prasad and J. S. Sahambi. Classification of ecg arrhythmias using multi-resolution analysis and neural networks. In *IEEE Convergent Tech*, pages 227–231, 2003.
- [19] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77, 2:257–286, February 1989.
- [20] J. L. M. Rita F. Fahy, Paul R. LeBlanc. Firefighter fatalities in the unites states - 2008. National Fire Protection Association Fire Analysis and Research Division, July 2009.
- [21] R. Silipo and C. Marchesi. Artificial neural networks for automatic ecg analysis. *IEEE Transactions on Signal Processing*, 1998.
- [22] M. R. Souryal, J. Geissbuehler, L. E. Miller, and N. Moayeri. Real-time deployment of multihop relays for range extension. In *ACM MobiSys*, 2007.
- [23] C. A. C. D. C. C. Stefanos N. Kales, Elpidoforos S. Soteriades. Emergency duties and deaths from heart disease among firefighters in the united states. *The New England Journal of Medicine*, 356(12):1207–1215, March 2007.
- [24] S. Yu and K. Chou. Selection of significant independent components for ecg beat classification. *Expert Systems with Applications*, 36(2):2088–2096, 2009.