

An Electronic Health Application for Disaster Recovery

Edmundo Chissungu
University of Cape Town
Room 300 CS Building
UCT, Rondebosch, RSA
+27 21 650 4752
chsedm001@uct.ac.za

Hanh Le
University of Cape Town
Room 314.10, CS Building
UCT, Rondebosch, RSA
+27 21 650 2667
hanh@cs.uct.ac.za

Edwin Blake
University of Cape Town
Room 312, CS Building
UCT, Rondebosch, RSA
+27 21 650 3661
edwin@cs.uct.ac.za

ABSTRACT

Natural disasters often destroy the fixed wired communications infrastructure. Therefore relying on such infrastructure after a catastrophe can be risky. Wireless Mesh Networks (WMNs) can provide wireless network coverage without relying on a wired backbone infrastructure or dedicated access points. In this paper, we focus on using WMNs in disaster recovery areas in order to disseminate information from patient to doctors. We investigate whether we can deploy an electronic-health application on top of an ad hoc peer to peer network. The initial prototype shows that it is feasible. However the application quality varies depending on the number of multi-hops and the applications concurrently running on top of the WMN networks.

Categories and Subject Descriptors

C.2.0 [COMPUTER-COMMUNICATION NETWORKS]:
General – *Data communications, Open Systems Interconnection reference model (OSI)*

General Terms

Measurement, Documentation, Performance, Design, Reliability, Experimentation, Human Factors.

Keywords

Disaster scenario, Wireless mesh networks, peer-to-peer networks, JXTA, E-health Application

1. INTRODUCTION

Natural disasters have a variety of causes and many different effects. However they inevitably impact the telecommunications infrastructure. When natural disasters, such as the 2010 earthquakes in Haiti, hit a populated area the infrastructure is damaged and often rendered completely unusable. The damage affects copper-telephone lines, fibre optic cables and Very Small Aperture Terminals (VSATs) in the area.

Disaster relief endeavours initially focus on rescuing as many people as possible. These rescue attempts are hindered by the lack of infrastructure to support the communications between the rescuers and those they are trying to rescue [1]. Relying on fixed terrestrial infrastructure has proven to be unreliable. This motivates the use of Wireless Mesh Networks (WMNs) which can provide wireless

network coverage of large areas without relying on a wired backbone infrastructure or dedicated access points (APs) [1-3].

We focus on using WMNs in areas where the communication infrastructure connectivity would have been destroyed as nodes will handle the routing of data amongst nodes that are still up (not destroyed by the natural disaster). In these areas WMNs can help send vital medical information from patient to doctors. We believe a decentralized and infrastructure free network to relay the appropriate medical information from patient to doctor is most appropriate. The necessary data would stream from a device on the patient, travel across an ad hoc network and end up at a sink (the doctor's computer) perhaps on the other side of the disaster zone.

In this paper we investigate whether we can deploy an electronic-health (e-health) application on top of an ad hoc peer to peer network. We measure the network's performance. Furthermore we report on the network's response to multiple applications running simultaneously on top of it.

In Section 2 the background to this research is covered. In Section 3 we describe the design of the network. We show the implementation of the WMN, and discuss the evaluation results in Section 4. The paper is concluded in Section 5.

2. Background

We start with a general overview of wireless mesh networks (WMNs) and their applications. The key components of our system are then presented: the WMN, the frontend for the patient, consisting of Electrocardiogram (ECG) and Accelerometer (ACC) sensor, and the backend for the physician, the ECG and ACC renderer.

2.1 Related Work

The idea of using WMNs as opposed to fixed wired infrastructure has existed for the last two decades. Portmann & Pirzada [1] provides an overview of WMN technology and its different uses in public safety and disaster recovery. Kumar et al [4], point out that one effective technology for disaster recovery is wireless ad-hoc networking. Rathy et al [4] focus on finding the most efficient routing algorithm for WMNs in a disaster recovery scenario: the requirement being a routing algorithm that guarantees that the packets sent on the network will reach its destination. They found that the dynamic source routing (DSR) algorithm is best suited since it has a low packet loss rate. However their work was a simulation study and we believe that in a real world setting coverage (the area the network reaches) would be most important. According to them DSR is especially suited in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SoICT'10, August 27–28, 2010, Hanoi, Vietnam.

Copyright 2010 ACM 978-1-4503-0105-3/10/08...\$10.00.

scenarios that have few hops. The problem with this is that most disaster scenarios span a very large area of affected people and so a WMN that can only cover a small area it is then not suitable. This is supported by Portmann, M., Pirzada, A., [5], who believe that WMNs can provide wireless network coverage of large areas without relying on fixed wired infrastructure or APs.

There are a number of ways in which WMNs can be used to service a disaster area: either by being deployed as a specific task or supporting different communication scenarios [1]. In our system we initially handle a specific task which is sending the medical information from the patient to the doctor across a multi-hop network. We then investigate if the network can support other applications running simultaneously on top of it.

A disaster WMN to service multiple disaster applications was proposed by Shibata et al [3]. They proposed creating a WMN in the sky above the disaster area using balloons. This approach is for those areas so severely damaged by the disaster that a new wireless network has to be set up before the recovery operation can begin. A prototype system was constructed to evaluate its function and performance through several disaster applications such as Voice over IP telephone (VoIP). Since this network can handle multiple applications it might benefit from the use of the peer-to-peer network such as ours constructed using JXTA to send patient information to their doctors.

2.2 Wireless Mesh Networks (WMNs)

The IEEE 802.11x family of standards (also known as WiFi) is currently the most successful broadband wireless networking standard for wireless local area networks (LANs) [6][7]. Nodes communicate without the need for an intermediate router, AP or gateway [8, 9]. Meaning that nodes in this network communicate directly with one another using each other as message router (devices in charge of delivering the message to its destination), and thus building an infrastructure-independent (independent in the sense that routers are not needed), multi-hop store-and-forward network [13].

3. System Design

In this section we discuss the details of the design that allow our peer-to-peer application to stream data from a patient, relay it across the network and deliver it to the doctors on the other side of the disaster zone. The system is divided into three main components namely Frontend-Patient, Backend-Doctor, and the core WMN as illustrated in Figure 1.

3.1 Frontend-Patient

The patient side is represented by the e-health application from which data stream. The data is generated by a sensor on the patient. The data is converted to an appropriate format and sent to the doctor's side. In previous work, we developed a Cardiac Monitoring System [10] which streams data from a ECG AliveTech [11] into a PDA. The application in the PDA reads data and converts that data to a readable format.

3.2 Backend-Doctor

The backend provides an interface for a doctor to view the patient's medical data as it arrives. We use our rendering engine that reads in the patients' electrocardiograph (ECG) and accelerometer (ACC) data and plots it on a graph [12]. The graph shows the patient's heart wave data as it would be seen on hospital equipment. It also plots the accelerometer (ACC) data which allows the doctor to know the patients physical orientation, whether standing, seated, or lying down. The

ACC data serves as a warning tool used for patients of advanced age that might have fallen and are unable to get up.

3.3 The Core WMN

The nodes will form a multihop network between nodes in the network which will allow for communication between nodes out of each other's radio range. This network will allow communication without the need for a router or access point (AP).

Figure 1 shows the network. Clouds represent the WMN network. The left-hand PDA represents the patient side with ECG devices. The right-hand PC represents the doctors receiving and displaying the data from the network

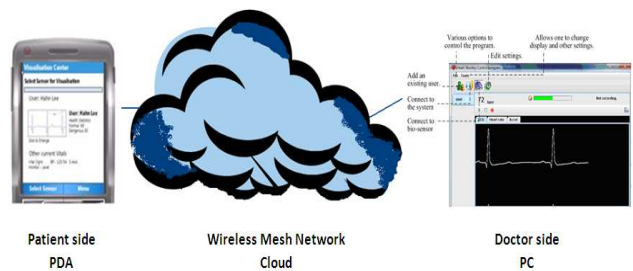


Figure 1. The System Overview

3.3.1 Detail of the WMN cloud

We used Juxtapose (JXTA), which is an open-source project defining a set of protocols for ad hoc and peer-to-peer computing irrespective of the underlying networking platform.

JXTA produces a virtual network on top of the physical network shown in Figure 2. At the very bottom layer of this virtual network are the interconnected devices like cell phones, PDAs, etc. In our case this would be the laptops or desktops with wireless adapters which joined the ad hoc network

The distinction between most networking packages such as the java.io package and JXTA's is that we do not use an Internet address or a port number; instead JXTA uses a user defined name for each node. This is an advantage as we are connecting personal computers (PC) that are not always addressable. This may be because a PC's address changes due to dynamic host configuration protocol (DHCP) or when laptops move between networks.

Each node behaves as a relay node for messages destined to a node with which it has no direct link level connection. There is also no need to have a server in order to have communication between nodes.

The JXTA sockets have an additional feature compared to customary sockets. There is an ID in the socket API called a pipe ID. The pipe ID acts as the socket's virtual address and port number. JXTA's routing system connects computers that are attempting to connect with the same pipe ID. For example, if a peer opens the JxtaServerSocket and listens for an ID of five, and another computer is looking for a pipe ID of five, the router connects the two.

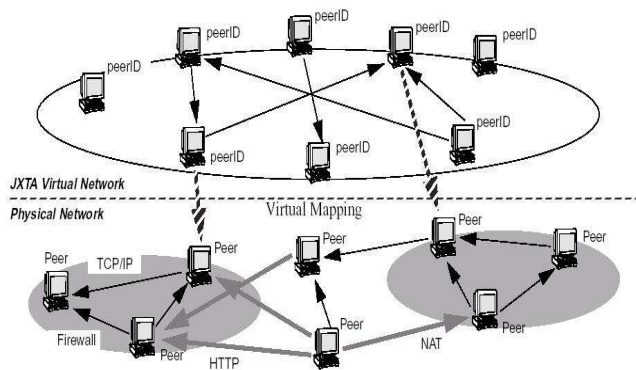


Figure 2. The JXTA virtual network

4. Implementation

The Electronic Health Application for Disaster Recovery enables doctors and patients who have been disconnected, due to destroyed infrastructure, to exchange necessary medical data. The physicians are then able analyse the data and monitor their patients caught in disaster stricken areas. We now describe how we set up and deployed the WMN on which the experiments were conducted. The technology chosen for the WMN was the ubiquitous IEEE 802.11x (WiFi) family of standards.

We used five desktop machines each running Windows XP professional edition service pack three. The machines were set up in a laboratory. We used WiFi cards and in each created a WMN profile called e-health_WMN and set the card to mesh (ad hoc or computer-to-computer) mode. We chose channel eight for communication.

We set up an initial wireless mesh network with fully connected nodes. In order to test the e-health application properly we needed a network with multiple hops between the patient node and the doctor node. With five nodes the maximum number of hops achievable is four: to achieve multiple hops in this small space we reduced the signal output strength and relied on the walls of the laboratory to attenuate the signals further.

We had to fine-tune the initial network topology to obtain a multi-hop network. Forcing multiple hops in a network set up in a relatively closed space was rather difficult to achieve. The difficulty is caused by amongst other things: omnidirectional antennae propagating signals throughout the space and the close proximity of the nodes.

While we did achieve four hops in the network (see Figure 3), the topology was unstable. This instability is due to nodes continually breaking the established network topology by communicating with nodes that originally they could not communicate with. For our simplified test bed a network with at least one to two hops suffices. Figure 4 shows the final network set up used to run the experiments.




Figure 3. WMN topology that gave us five hops

In Figure 4, the left PDA represents the patient ECG devices. The right PC represents the doctors receiving data from the patients. There are two hops in the testbed to connect the backend and the frontend systems.

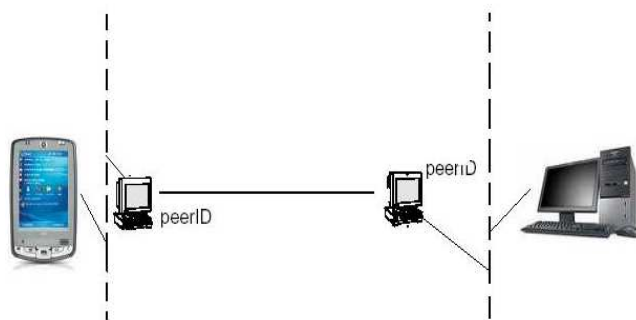


Figure 4: The final system prototype

4.1 Evaluation

We first explain how the experiments were conducted and then present our results.

The experiment consisted of having a patient and doctor node interact with each other. The patient node sent a file containing ECG and ACC data. The file had 2598 seconds (43.3 minutes) of recorded data: 1368921 bytes. We chose to send segments worth a second of application data through the network, which is roughly 525 to 527 bytes of which 300 bytes were ECG and 225 bytes were ACC data. This resulted in 2598 messages in total.

Table 1. Average time to send complete file

Link Number	Average time in milliseconds
1	938.36
2	908.063
3	1030.812
4	634.281

We measured the traffic at each link on the network topology in Figure 3. We had 2598 data points of raw data. Table 1 shows the average time it took for the complete file to be sent over a particular link. The link number, n, represents a link shared between nodes n and n+1.

Table 2 shows the average throughput of each link (bytes/second).

Table 2. The through for each link in bytes per second

Link Number	Throughput (bytes/seconds)
1	1441
2	1504
3	1325
4	2153

Tables 1 and 2 shows that the individual links can support the deployment of the e-health application. From Table 1 we can see that the file with 2598 seconds of data traversed the slowest link (link 3) in its entirety in 1030.812 seconds. This transmission time is significantly lower than the time represented by the recorded data stored in the file. This suggests that the e-health application data can be transferred in real time. This is further supported by the data presented in Table 2, where the throughput is higher than the byte blocks of patient data sent at any given time. Link 3, once again, has the lowest throughput, with a value of 1325 bytes per second. This confirms that we can deploy an e-health application on top of an ad hoc peer-to-peer network since the link speeds, even with the fluctuations seen in Table 2, are sufficiently high.

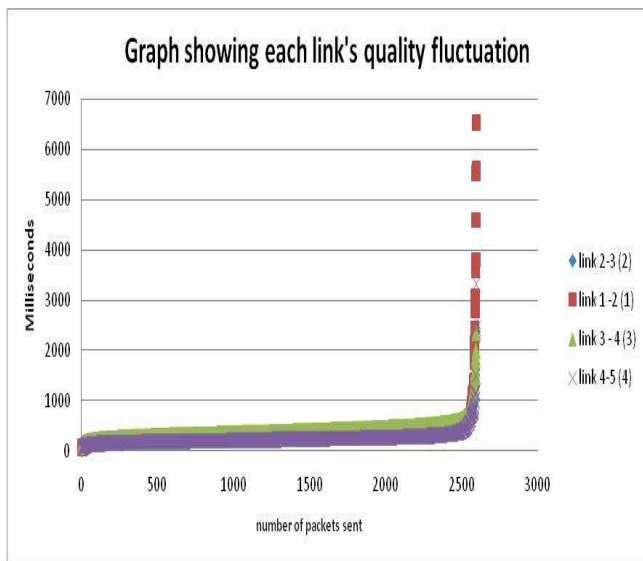


Figure 5. Shows the link fluctuation for each of the links by showing how long it took to send each packet/byte block

Figure 5 plots the links' quality throughout the duration of the experiment. The Y-axis represents the time in milliseconds and the

X-axis represents the number of packets/byte blocks sent. As can be seen in Figure 5, the link quality of the WMN is fluctuating. Link one and link three shows the two extreme values in this fluctuation.

We wished to test the networks response to multiple applications running simultaneously on top of it. We expected to see the e-health applications throughput to decrease but did not know by how much. We performed the experiment with another application, similar to the e-health application, running on the network. The e-health application throughput was 640 bytes per second. This value is low, as expected, however it is still sufficient to send the data without too much delay. The e-health application under the scenario mentioned above managed to send the file containing the patient's data in its entirety in 2131.563 which is lower than the amount recorded in the file. However, this data suggests that if a third application were to run simultaneously on the network then the throughput would drop significantly and the time to send the patient's data would take longer than the time the data represents.

5. Conclusion

We investigated whether the deployment of the e-health application on a peer-to-peer network was feasible. Our findings show that the e-health application can be deployed on a peer-to-peer network. The intended use of the application is in a disaster recovery scenario where existing networks have broken down and have to be replaced by emergency ad hoc connections. The application makes use of the most suitable infrastructure for disaster recovery operations, namely Wireless Mesh Networks. A typical test situation was recreated where realistic ECG data was successfully sent across the network.

However, our findings also revealed a problem when the network has more than two applications running simultaneously. This is especially true if the applications are similar to the e-health application, as both require similar amounts of resources. One possible way to mitigate this might be to implement a queuing system where the applications with higher priority get sent first. The priority can be determined by the importance of the application to its users. In a situation where this approach is too simple other quality of service assurance skims will have to be investigated.

Our papers' main contribution is to stir in our readers an interest in the use of technology to better serve people in impoverish areas whether by a natural disaster or social economic factors. This is in an endeavour to help our fellow man through the correct application of technology.

6. Future work

We hope to see whether the live streaming of medical information can be supported by the peer-to-peer network. In this scenario we would have a patient continuously sending medical data across the network and the doctor streaming this data and displaying it as comes in. The effect will be of the doctor monitoring the patient's vitals as if they were physically present in the room.

7. ACKNOWLEDGMENTS

Our thanks to Nina Schiff, a computer science student at the University of Cape Town (UCT), for writing our rendering engine.

8. REFERENCES

[1] M. Portmann and A. A. Pirzada, "Wireless Mesh Networks for Public Safety and DisasterRecovery Communications," A. s. I.

- R. C. o. Excellence.nicta.com.au, Ed. Australia: Australia's ICT Research Centre of Excellence.
- [2] D. Windle, "The Role of Wireless Technology in Disaster Recovery," KingsBridge Disaster Recovery 2002.
- [3] Y. Shibata, Y. Sato, N. Ogasawara, and G. Chiba, "Ballooned Wireless Mesh Network for Emergency Information System" presented at International Conference on Advanced Information Networking and Applications - Workshops (aina workshops 2008), March 2008.
- [4] K. Suresh, R. K. Rathy, and D. Pandey, "Design of an ad-hoc network model for disaster recovery scenario using various routing protocols," in Proceedings of the International Conference on Advances in Computing, Communication and Control. Mumbai, India: ACM, 2009.
- [5] D. Johnson and A. Lysko, "Overview of the Meraka wireless grid test bed for evaluation of ad-hoc routing protocols," in Southern African Telecommunications and Networks Conference, September 20.
- [6] M. J. LEE, J. ZHENG, Y.-B. KO, and D. M. SHRESTHA, "Emerging Standards for Wireless Mesh Technology," in IEEE Wireless Communications Magazine, vol. 13 2, April 2006.
- [7] C. VINCENT, "Understanding link quality in 802.11 mobile ad hoc networks," 2004.
- [8] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey, Computer Networks," vol. 47: Elsevier Science Publishers B. V., 15 March 2005, pp. 445-487.
- [9] T. Huovila, P. Lassila, J. Manner, and A. Penttinen, "State of the Art Analysis of Wireless Mesh Technologies," November 2006.
- [10] L. Hanh, K. Michel, and C. Gregory, "An Electronic Health Care - Cardiac Monitoring System," presented at ICC 2010 Workshop MAN, Cape Town, South Africa, May 2010.
- [11] AliveTech.
- [12] H. Le, N. Schiff, J. d. Plessis, and D. Hoang, "A Pervasive Telehealth System for continual and low intrusive monitoring using Peer-to-Peer Networks," presented at In Proceedings 5th International Conference on Information Technology and Applications, Cairns Australia, 2008.
- [13] S. Asherson and A. Hutchison, "Secure Routing for Wireless Mesh Networks," presented at In Proceedings Southern African Telecommunication Networks and Applications Conference, Spier Wine Estate, Stellenbosch, South Africa, 2006