Wireless communication in telemedicine using Bluetooth and IEEE 802.11b

Magnus Berggren

Department of Information Technology
Uppsala University
Box 337, SE-751 05 Uppsala, Sweden

Technical report 2001-028
November 2001
ISSN 1404-3203

Uppsala Master's Thesis in Computing Science
Examensarbete DV3, 1DT150
2001-11-16

Wireless communication in telemedicine using Bluetooth and IEEE 802.11b

Magnus Berggren

Information Technology
Department of Computer Systems
Uppsala University
Box 311
SE-751 05 Uppsala
Sweden

Abstract

This thesis explores the issues surrounding the simultaneous deployment of Bluetooth & Wireless LAN (IEEE 802.11b) for networked devices in a telemedicine environment. A key aspect in telemedicine is the ability of the two standards to co-exist in close proximity. This work defines telemedicine, gives examples of applications, and describes the two wireless techniques and the issues that arise when considering the introduction of wireless communication.

The thesis tries to answer the question whether the two wireless techniques Bluetooth and Wireless LAN can co-exist in a telemedicine environment. In order to do this, quantified measurements where performed were the two wireless techniques where exposed to radio interference from the other technique. These results are printed as graphs and explained.

Supervisor: Professor Per Gunningberg
Examiner: Professor Per Gunningberg

Table of contents

1 Thesis specification / Problem formulation 4
2 Telemedicine 4
2.1 Definition 4
2.2 History of telemedicine 5
2.3 The telemedicine project 5
2.4 General Objectives 5
2.4.1 Technical demands 6
2.5 Agilent 6
2.6 Ortivus 7
3 Wireless technologies 9
3.1 Health issues 9
3.2 Bluetooth 10
3.2.1 Protocols 11
3.2.2 Piconets 11
3.2.3 PCMCIA-card from IBM 11
3.3 IEEE 802.11b, Wireless LAN 13
3.3.1 Wi-Fi 14
3.3.2 PCMCIA-card from Lucent 14
3.4 Reference data rates of the two techniques 15
3.4.1 Bluetooth 15
3.4.2 Wireless LAN 15
4 Experimental Variables 15
4.1 Measurement metrics 16
4.1.1 dBm 16
4.1.2 Signal-to-Noise Ratio (SNR) 16
4.2 Problem with signal propagation 16
4.2.1 Problems 17
4.3 Distance vs. lost packets 17
4.4 Influence of movement 17
4.5 Interference from competing radio techniques 18
4.6 Method 18
4.7 Test environment 18
4.7.1 Computers 18
4.7.2 Test locations 19
5 Interpretation of results 20
5.1 Bluetooth interference 21
5.1.1 Bluetooth disturbing WiFi 22
5.1.2 WiFi interfering with Bluetooth 23
6 Conclusion 24
7 References 25
8 Appendix 28
8.1 A, Sniffer programs 28
8.1.1 Commview 2.4 28
8.1.2 NAI Sniffer 28
8.1.3 Analyzer, WinDump and WinPcap 28
8.1.4 TTCP, PCA_TTCP and NTTCP 29
8.1.5 Iperf 29
8.1.6 Agilent Analyzer, J1955A 29

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

8.1.7 WS_ping ProPack 29
8.2 B, The client-server program 30
8.2.1 The client 30
8.2.2 The server 30
8.3 The timetable 31

1 Thesis specification / Problem formulation

Neonatal infants often have to stay in hospitals for long time periods, sometimes up to several years, because their condition demands qualified supervision. These long periods of hospitalization have a negative influence on both the parents and the child. In addition this form of care is costly.

There is an extensive interest from the medical service, the educational institutions, the industry and the parents to review how the present and future wireless communication techniques will influence telemedicine, especially home environments. Among the families who are taking care of their child at home there is a need for increased mobility and a desire to have easier management of the needed equipment. The medical institutions might lower their costs by adopting newer techniques as well as enriching the life of the care-needing families. The academic institution is interested in researching how these techniques function in new applications and how the demands differ. The purpose of this thesis is to examine if wireless techniques like Bluetooth and Wireless LAN can provide solutions that reduce the number of cables and wires used.

One of the problems with using radio in this context is the possibility of interference from other radio sources such as microwave ovens, cordless phones and competing networking products using the same frequency band like Wireless LAN and Bluetooth.

When two or more devices using the same radio technique communicate they share the medium between them in a way that minimizes the possibility that they send data onto the ether at the same time. If other devices using another radio technique are in close proximity the formers might hear their communication but since the radio is used in another way the formers can not interpret the transmission, thus it is considered as noise. High levels of noise might cause communicating devices to lose connection or cause corruption of the transferred data.

When you know that radio connections suffers from these and more problem, questions arise, such as how does the connection fail, what are acceptable levels of noise, how low signal strength should we accept, what is the right thing to do when the connection is lost. In the wireless section of this document some of these questions will be explained in more detail and answered.

This thesis is a part of a larger project that is a collaboration between the Academic children's Hospital (ACH) and the Department of Computer Systems (DoCS) at Uppsala University. DoCS has the task of supporting ACH with the technical aspects of the project.

The work consists of reading and learning about telemedicine, familiarizing with telemedicine terms and products, learning Linux, programming of a client-server network-measuring software and reading and learning about Bluetooth.

2 Telemedicine

> *Enligt Världshälsoorganisationen (WHO) är telemedicin: "...utövande av hälso- och sjukvård genom användande av interaktiv kommunikation av ljud, bild och data. Detta inkluderar diagnossättning, konsultation och behandling* såväl som utbildning och överföring av medicinska data".

2.1 Definition

Telemedicine is a generic term referring to all forms of medical information exchange, including a variety of telecommunication technologies. Applications in health- and medical care include

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

telecommunication, data- and information technique are used to transfer medical information. The way in which the information is transferred depends on the situation, technology and conditions such as planned or spontaneous, synchronous or asynchronous.

The purpose of telemedicine is to provide medical expertise to places that due to their physical distance do not have access to this kind of competence. Examples of environments where these techniques could be of use could be as decision support for pre-hospital personnel working in ambulances or in patients' home. Telemedicine are also used to transfer important medical information from the remote location to a hospital or some other institution. The use of telemedicine will also make it possible to create organizational and structural changes such as out-of-hour working, homecare and journal writing in different time zones.

2.2 History of telemedicine

The development of telemedicine has been growing during the last 20 years. However, it is during the last five years that the development has had its largest growth. A search in the database MEDLINE [Medline] on the word "telemedicine" gave 296 hits between the year 1966/01/01 and 1994/12/31. A search between 1995/01/01 and 2001/05/10 resulted in 3780 hits, of which 785 were between 2000/01/01 and 2001/05/10.

Developments in telemedicine was from the beginning driven by the military, space research, sea fare- and aviation industry and was mostly concentrated in remedy alarm situations were the geographical distances made it impossible to transfer injured persons to hospitals or were qualified personnel are not available. Recent developments are increasingly focused on problems were the geographical distance is not the main focus, instead the focus has been cost reduction and inconvenience of traveling.

The first tests of telemedicine in Sweden took place in 1922 when the Sahlgrenska Hospital started to give medical advise to ill and injured sailors via radio. The first conventional tests of telemedicine were conducted in Uppsala 1968 within neurophysiology.

The Swedish telemedicine practice expanded rapidly in the middle of the 1980's in the wake of the pc revolution, and it has been expanding at an even faster rate since the mid-1990's when Internet became publicly available. At the beginning, telemedicine was mostly used for teleradiology, which means transferring X-ray pictures. With the increase in bandwidth and faster computers, new applications such as videoconferencing systems were developed. As the technology is getting more mature the use of telemedicine is expected to be more and more an integrated part of the daily work in the hospitals and homecare.

The project Sjunet could serve as an example. Sjunet is a computer network built by seven county councils in the Uppsala-Örebro region. The network is a part of the Telia backbone as a Virtual Private Network (VPN) with full integrity and security. Sjunet offers e-mail, videoconferencing, EDI and connectivity to the Internet.

2.3 The telemedicine project

The telemedicine project consists of two active participants, the Department of Computer systems at Uppsala University (DoCS) and the Academic children's hospital in Uppsala (ACH). The project is founded by a grant from the Uppsala county council. The hospitals role is to work out new routines and forms, which allows the introduction of telemedicine in the daily work. DoCS has the role of evaluating different technical aspects that have to be taken into account when using telemedicine in a home environment.

2.4 General Objectives

What we want is a system that can be placed in a home environment where the demands differ from those on a hospital. To begin with, we want a system that is small and mobile which enables the

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

parents to take the child and the system with them as they move around between different rooms in the household. As a result of this we would like the system to be able to run on batteries for some time thus allowing the system to be on when for example the parents take their child to the hospital.

The user interface has to be designed so that it is easy to understand. In a critical situation parents have to be able to understand what is happening in order to make the right decision. When an alarm is triggered they have to be able to make a quick and easy decision if they have to go to the hospital, if the alarm is an artifact or if it is some kind of malfunctioning hardware. They want to know the severity of the alarm that goes off, or in the case of multiple alarms, they want to know which of the alarms that is the most severe. If a cable attached to the patient detaches, an alarm that indicates this should go off. If, for example, a cable to a sensor that measures heart rate detaches the patients pulse will drop to zero indicating a stopped heart, which in turn will cause a more severe alarm to go off. In this situation it could be fatal if the parents start doing heart massage instead of re-connecting the cable. The system has to be designed so that these situations are avoided.

Not every home has a permanent Internet connection. Since the ordinary telephone line might be occupied with a phone call we have to take into account other solutions such as GSM, Mobitex, or in the near future 3G solutions. Wireless broadband wide area coverage will certainly give the telemedicine revolution even more speed, enabling system the possibility of being constantly connected, higher data speeds, mobile and yet just paying for the amount of transferred data and not on a time basis as with GSM, which mean that the equipment could be constantly connected.

This leads us to the last aspect of the demands on the system, wireless communication. There are numerous aspects that have to be taken into account when you change from a wired to a wireless network. Problems with interference from other radio sources transmitting at the same wavelength, losses due to distance between sender and receiver, incompatible protocol implementations, increased power consumption, unhealthy radiation from the radio transceiver, higher prize, security issues and possibly lower bandwidth. Some of these issues will be addressed later. But besides from all of these problems the real benefits are no cables that can be “detach” or stumbled upon, and the increased mobility.

2.4.1 Technical demands

The system should be a small and reliable monitor. It should have sensors that can collect ECG for determining of RR-interval (frequency), thorax impedance for determining of breathing motion, frequency / apnea, sensors for pulsoximety, which gives saturation, and pulse frequency. The sensors have to be fixable so that continuous monitoring can be guaranteed. The signals should be filtered in such way that it sorts out artifacts. The alarm levels should be adjustable and smart. The primary goal is to register the variance in the RR-interval of the ECG, which indicates decreasing heart rate. If the heart rate drops below 80/minute for more than 5-10 seconds an alarm should go off. An algorithm should then check if a bradykardi breath motion exists and if the pulsoximetry measurement indicates saturation levels over 85% with adequate pulse amplitudes.

The monitor should store frequency data as trends, for example as histograms, with information about the number of alarms and events outside of the set alarm levels. Information should be presented on a display to the parents in the home, but it should also be possible to send the information to the hospital or to a moving car. If an alarm is not signed for within a certain amount of time, the alarm should be sent to the hospital along with additional monitor information. If technically possible a voice and picture channel would be an interesting feature as well.

2.5 Agilent

The Agilent Viridia M3 monitor, picture to the right, is a small lightweight patient monitor for lower acuity care, emergency and transport applications for both adults and neonatal patients. The monitor was in October 2000 equipped with a built-in wireless LAN from Proxim Inc [Agilent2000]. This information is transferred to the Agilent Information Center. Switching from

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

wired to wireless network is done by just removing the network cable. The monitor has a four-color display with a touch screen interface capable of displaying up to four waveforms simultaneously. The monitor has the ability to measure ECG, Respiration, SpO_2 and non-invasive blood pressure (NBP). It has also a connector for either temperature or invasive blood pressure and optionally CO_2.

Agilent Viridia monitor

Agilent Telemon Monitor

The most recent monitor released by Agilent is the Telemon, picture to the left. The Telemon is their smallest monitor and it has the ability to measure NBP. The monitor is equipped with the Agilent Telemetry System that measures ECG and SpO2. The Telemetry System has a built in Wireless LAN card which sends the measured data to the Agilent Information Center for further analysis. The choice of Wireless LAN can, as we will see later, cause implications for devices using Bluetooth.

Our initial attempt was to take the Agilent Viridia monitor and extract data from it and transfer the data on some kind of network to a server that could present the data onto a computer screen. This was however not possible in the summer of 2000, the only data that could be extracted from the monitor was a screen dump. This led us to look for an alternative to the Agilent solution.

2.6 Ortivus

After some investigations we found that Ortivus was developing a system that could fill our needs after some modifications and additions. The system consists of three parts, a measurement unit (MU), a presentation/analysis unit (PWS) and a server unit. The system is based on the Mobimed system, which is used by nearly 50% of the Swedish ambulances, which are online with some 30 hospitals across the country [Mobimed].

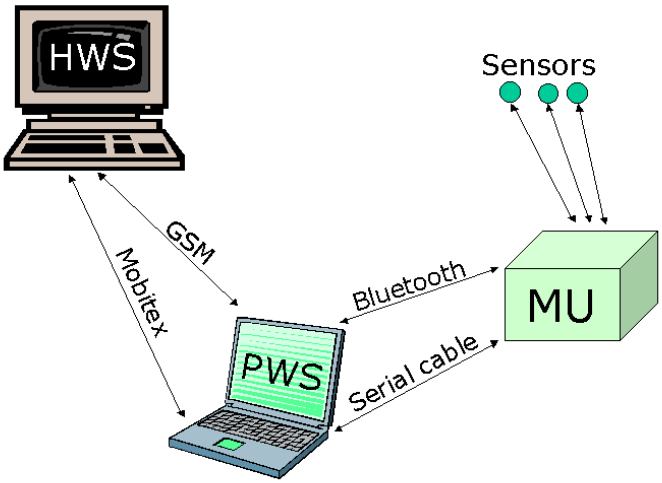

Picture: The Ortivus system

The MU is, at its present shape, a small box the size of two VHS-cassettes on top of each other. This unit has the ability to measure VCG with derived 12-lead ECG, oxygen saturation SpO2 and non-invasive blood pressure (NBP). This little device has a small 8-bit processor that works with a load of nearly 90% of its capacity with just sampling the data and compressing it. There is no spare processing capabilities for doing analysis etc. The device has also a small memory. This memory has the capacity of storing a few seconds of sample data. This doesn't give any room for re-transmitting corrupted data or storing data while waiting for re-establishing of a broken link.

This device collects and routes the gathered information to the presentation and analysis unit via a serial interface or, as in Ortivus' research lab, via a Bluetooth transmitter. This issue will be addressed later in this section.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

The presentation and analysis unit, PWS, is a portable laptop with or without touch-screen functionality running under Windows 2000. This device receives the data sent by the MU either through the serial cable or via the Bluetooth receiver. It is then the job of this device to analyze the data and display the result on the screen. Using a standard portable computer with Windows 2000 for the analysis and presentation allows software to use integrated components such as libraries, drivers, high-level languages and compilers. Using standard equipment has the advantage that a smaller portion of the equipment has to go through the thorough medical certification procedures. This part of the system sets off alarms if it detects a dangerous situation. The system can trigger several different types of alarms. If the data connection between the MU and the PWS is lost one type of alarm is triggered. This alarm is not critical if the patient otherwise is fine. Another type of alarm, called artifact, is triggered if one of the measuring devices attached to the patient's body detaches. For example, if one of the diodes measuring heart rate detaches causing the measured heart rate would drop to zero. A less severe alarm should be triggered, informing the monitoring person to reattach the loose diode. And of course, if a severe alarm is triggered, possibly caused by a stopped heart or stopped breath motion, a critical alarm must be triggered.

Future developments will be to combine all the gathered information in order to find patterns that might lead to severe alarm situations. For example, if the frequency of the breathing is varying too much this might lead to something more critical if it is allowed to continue. If it is possible to prevent this from happening an alarm should go off, preceding the dangerous situation.

The server system is placed at the hospital. The (PWS) workstation placed in the homes of the patients sends information to the server via the wireless interfaces Mobitex or GSM. However it is possible to use wired networks if that is wanted. The patient workstation connects to the server on a regular basis and sends information about the current situation. If an alarm goes off, the PWS automatically connects to the server and starts sending real-time data. This information can be displayed on the server's monitor by the hospital staff, allowing the hospital staff to guide the parents to a solution to the problem, send an ambulance or instructing them to come to the hospital for further treatment. If the child has to be moved to the hospital the monitor could still be attached and sending live data to the server during the travel. If the connection gets disrupted, e.g. due to driving through a tunnel, the PWS stores the information and resends the data as soon as the link comes up again. Due to the low bandwidth of the GSM and Mobitex networks the information has to be compressed, but even so the communication device has to reserve space in order to be able to resend data.

Separating the gathering unit from the presentation and analysis unit is a new approach. There are several reasons for this separation. The most important is that medical equipment has to go through time-consuming certification procedures to get licensed. By separating the parts that requires medical approval from parts that don't need it you can later change the specifications for parts that aren't covered by the license without having to go through the certification process again. This means that you can take advantage of the rapid development of computer hardware and upgrade parts that are not covered by the medical license. The second reason has to do with mobility. By reducing the size and weight of the unit that has to be attached with wires to the body you increase the mobility. The third reason is since the PWS device has a lot more processing and storage capabilities than the MU, it is capable of doing more complex tasks that wouldn't be possible on the MU and it is also capable of storing the collected data for a longer time.

The advantage of using a standard Windows 2000™ operating system on the PWS and server is that applications can be developed in high-level languages with a vast variety of libraries, classes and drivers available. Windows systems are, in general, more familiar than those based on Linux or an embedded operating system.

A negative aspect is that Windows lacks support for real-time processing. With a Windows machine there is no guarantee that a task will be finished within a given time. There might be other tasks that the operating system prioritizes more. Another disadvantage is that Windows is known to

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

be less reliable than other operating systems. The total cost of ownership (TCO) is also higher for Windows than for Linux. The license cost for example is higher on Windows, but this might not be a big deal in this context where other costs are much higher. The security is lower on a Windows system than on a Linux system. Windows systems are known to be an easy target for those with malicious intentions. And last, a Windows machine consumes much more resources, in terms of battery consumption and storage needed, than a Linux or embedded system. This will decrease the mobility aspect and increase the unreliability of the system since the battery time will be shorter.

If Ortivus' tests with a Bluetooth connection between the MU and the PWS work out well, the Bluetooth functionality of piconets (explained below) can be exploited. In the future, the MU can be splited into several smaller components where each of these has just one specific tasks only measuring one thing each. These smaller parts can then be combined like Lego™ resulting in a more specific system with less or the same functionality as the present monitoring unit. The benefit is that Ortivus don't need to make a special device for every individual customer. A device will not have more hardware than the customer has paid for, resulting in a less expensive, smaller, easier and cost-effective systems. What makes this possible is the functionality offered by Bluetooth piconets. The piconet will tie these parts together. Bluetooth piconets allows up to eight nodes where seven of them act as slaves and one is the master. If Ortivus carry through this idea with Bluetooth communication the PWS will act as master and up to seven small slave devices can be connected to the master.

3 Wireless technologies

The two solutions, Wireless LAN and Bluetooth, offer wireless networking capabilities with different focus. Bluetooth focuses mainly on offering a solution that reduces the amount of short-range cables required. In the telemedicine context the aim would be to find solutions that eliminate the cables from the sensors to the measurement unit or from the measurement to the analysis unit, where the devices are located near each other, probably in the same room. The focus of wireless LAN is to offer networking capabilities like a wired LAN with higher bandwidth and longer range than Bluetooth. A Wireless LAN solution offers connectivity through Access Points connected to a wired network. A WLAN will cover several rooms, floors or possibly a whole building. This makes it possible to monitor a patient and send the data to the central unit while moving the patient from one room to another or to some other part of a building.

So in terms of capabilities and focus these techniques don't compete, they rather complement each other, but in terms of frequency usage the answer is not that obvious. They share the same frequency band but can they co-exist?

The following section provides an overview of the two techniques Bluetooth™ and Wireless LAN a.k.a IEEE 802.11b a.k.a Wi-Fi™. Both of these uses the 2.4 GHz globally license free Industrial, Scientific and Medical (ISM) band and comply with the regulations for this band.

3.1 Health issues

Bluetooth and Wireless LAN use the same frequency range as microwave ovens. We all know that the microwave oven heats food that contains water. The obvious question that follows is then "What about the dangers of getting in the way of such radio waves?"

The transmitting power of these devices is by far too weak to be noticeable to humans. Also, the radio antenna sends omni-directional and not as a beam. But, if you use a mobile phone, headset or similar product equipped with either of these radio techniques some of the microwave radiation energy will hit the body. When using the 2.4 GHz band the penetration depth into a body is no more than 1.5 centimeters, which means that the absorption is superficial. The main absorption mechanism is field-induced rotation of polar molecules that generates heat through molecular friction. Water has its natural molecular oscillation at 2450 MHz, which is in the middle of the ISM

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

band. This feature is used in microwave ovens to make the water in food cook. However, a normal microwave oven uses power levels one million times greater than the power used by most of these devices. Research has shown that the output power of these communication devices is far too low to cause any detectable temperature increase of the human body.

The IEEE, WHO and ICNIRP have developed Radio Frequency exposure recommendations specifying near-field restrictions referred to as SAR between 10 MHz and 10 GHz. High power Bluetooth devices might need to be tested but the low power devices wouldn't.

Recent research performed by Lennart Hardell at the regional hospital in Örebro (RSÖ) has shown that the use of mobile NMT mobile phones increases the chance of getting cancer in the temple lobe. This means that there has to be some awareness to these issues when dealing with radio devices.

3.2 Bluetooth

Bluetooth is a low cost, low power, short-range radio technique introduced by Ericsson and others. Bluetooth is originally and essentially a replacement for physical cables. The goal of eliminating cables has lead to the creation of the notion of Personal Area Networks (PAN), a close range network surrounding a person carrying several heterogeneous devices equipped with wireless communication techniques.

The aim is to make connectivity simple, seamless and intuitive. Users should no longer have to install, plug in, enable or configure their devices in order to get them to communicate with each other. With the fulfillment of this goal the communication will be seamless and ubiquitous to the user. However, this simplicity to the user puts the demands on the Bluetooth developer who has to deal with all the complexity.

Bluetooth devices operate at 2.4 GHz in the license free Industrial Scientific and Medical (ISM) band. This band specifies a basic set of power, spectral emission and interference rules. In order to be a winner in this shared frequency band Bluetooth has focused on being robust at the cost of transfer speed. The ISM band ranging from 2.4465 to 2.4835 GHz is divided into twenty-three 1MHz channels, each signaling data at maximum of 1 Mega-symbol per second. Bluetooth sends one packet in one of these 1 MHz slots and then both the receiver and the sender tunes into another channel. This results in a hopping process that eventually will make sure that all of the 79 channels will be used. This behavior is called FHSS – Frequency Hopping Spread Spectrum. Bluetooth hops between channels 1600 times per second. This means that every timeslots is 625 microseconds, which is the normal time it takes to send a packet. There is however the possibility of sending larger packets spanning over 3 or 5 timeslots resulting in increased transfer speeds, see table below. The reason for this fast hopping and short packages is robustness. If a part of the ISM band is exposed to interference and thus compromises the transmission other parts might still be available and the re-transmission will hopefully be sent in a frequency where interference is not present.

In order to be low power the Bluetooth specification specifies three different power modes; 1mW (0 dBm) which is the standard level, 2.5mW (+4 dBm) and 100mW (+20 dBm). These modes will give the device an operational range of 10, 20 or 100 meters.

Bluetooth supports two kinds of links, Asynchronous Connection-Less (ACL) and Synchronous Connection-Oriented (SCO). ACL provides packet-switched connections between the master and all active slaves. Packet retransmissions are usually applied to ensure data integrity. SCO links are usually used for voice transmissions. The SCO link is point-to-point between the master and one slave. The master maintains the link by using reserved timeslots at regular intervals to poll the slave. In SCO, packet retransmissions are not allowed. Bluetooth supports one ACL channel, three SCO channels or one channel that simultaneously supports asynchronous data and synchronous voice.

Bluetooth uses the network in an Ad-hoc fashion. This means that a network is formed between multiple Bluetooth devices in a given coverage area without the use of an access-point or a

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

server. In Bluetooth all communication goes through the master unit. There is no direct communication between the slaves. The idea is not that the master should route information between the slaves. Instead, if two slaves want to communicate with each other they can form a new piconet with one of them acting as the Master. However a device is not forced to leave the previous piconet, Bluetooth specifies a mode in which devices can be set as parked in the "old" piconet while they communicate in the new. Fast reconfiguration is one of the key features in Bluetooth.

3.2.1 Protocols

Bluetooth consists in general of three core protocols: The logical link control and adaptation protocol (L2CAP), the service discovery protocol (SDP) and the RFCOMM protocol.

L2CAP, which adapts upper layer protocols over the Baseband, provides data services to the higher layer protocols with protocol multiplexing capability, segmentation and reassembly operations and group abstractions. Device information, services and the characteristics of the services can be queried using the SDP.

Like SDP, RFCOMM is layered on top of the L2CAP. As a cable replacement protocol, RFCOMM provides transport capabilities for high-level services that use serial line as the transport mechanism.

3.2.2 Piconets

A piconet is essentially a collection of devices connected via Bluetooth technology in an ad-hoc fashion. A piconet generally starts with two connected devices, such as a portable PC and a mobile phone. The limit is set at 8 units in the piconet (that is why the required address-space is limited to 3 bits). All Bluetooth devices are peer units and have identical interfaces to each other. However, when establishing a piconet, the initiating unit will act as a master for synchronization purposes. The master's clock and hopping sequence are used for this synchronization. The other unit(s) will be slave(s) for the duration of the piconet's existence.

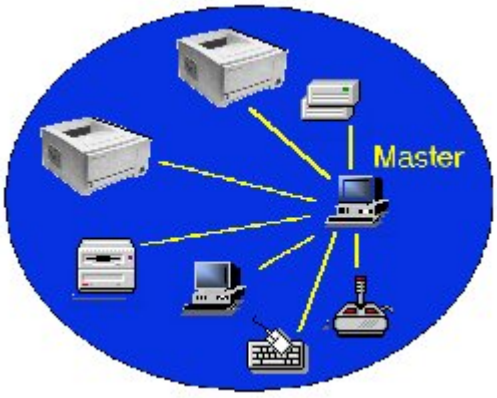

The extension to piconets is called **scatternets**. This is when two or more independent and non-synchronized piconets communicate with each other. A slave or a master in one piconet can become a slave in another piconet. If the need arises this unit can then relay information between the two piconets.

A unit cannot be a master in two piconets. If a master in one piconet wants to start yet another piconet and thus become the master of that piconet as well, Bluetooth provides primitives that makes it possible to change master in a piconet during operation, allowing mastery to be delegated to another node.

3.2.3 PCMCIA-card from IBM

In the following tests PC-cards from IBM with product number '09N9812' where used. These cards are sold with drivers with software from Digianswer.

3.2.3.1 Technical specification

These cards operate in the 2.4 GHz ISM unlicensed radio spectrum, they support a maximum data transfer rate of 724 Kbits at a range of up to 10 meters, secure data and voice communication and have low power consumption. The cards are compatible with the Bluetooth 1.0B standard and have upgradeable firmware.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

The cards have support for Bluetooth standard transmit power modes, have an integrated diversity antenna and are delivered with a user application software which includes profile management and control center.
The cards support the following profiles:

- Generic Access Profile
- Service Discovery Application Profile
- Serial Port Profile
- Headset Profile
- Dial-up Networking Profile (as Data Terminal)
- FAX Profile (as Data Terminal)
- LAN Access Profile (as Data Terminal)
- File Transfer Profile
- Synchronization Profile (supported by the API)

3.2.3.2 Availability and price

Bluetooth PCMCIA-cards are available on the market since October 2000 from manufactures such as IBM, Toshiba and Motorola, all of them produced by Digianswer. The price for an IBM card is $189 + taxes and shipping.

3.2.3.3 Problems with the driver

Problems have been experienced during the project with the behavior of the Digianswer drivers. The behavior is not the same in Windows 98 as in Windows 2000. If Windows 98 is used as sender of data and the Windows 2000 as receiver, the connection looses packets. If the computers are sending in the opposite direction the behavior is different. Then the connection will not drop any packets, instead the throughput decrease as the connection get weaker. The probable reason for this is that the Windows 2000 driver is implementing secure transfer at the baseband and Windows 98 is not. The driver can specify an option called Flush Timeout [BT 1.1, section 6.2]. A value of 1 units of milliseconds implies no retransmissions at the Baseband level since the minimum polling interval is 1.25 ms. If the value is all 1's this means that there should be an infinite amount of retransmissions, this is also referred to as "reliable channel". This means that the link manager should continue to send a segment until the physical link is lost. This value is asymmetric, the sender of a Request should specify its Flush Timeout value if it differs from the default value of 0xFFFF.

This means that the two implementations of the Bluetooth driver from Digianswer can differ on this point. This would explain the different behavior of Windows 98 and Windows 2000. I have tried to contact Digianswer about this without getting any answer.

According to the Digianswer SDK primitives are provided to read this value when writing programs for the driver. However since this is not in the scope of this work such an investigation has not been conducted.

The question following this observation is "Will this affect the measurements?" The probable answer is no. In order to explain this some words has to be said about the two protocols, UDP and TCP. UDP provides unreliable packet based transfer. The normal usage includes for example streming media or networked games, for example if you listen to music over the web it is not a big deal if you loose some of the packets, the resulting gap will not be heard. You would not want to wait for a lost packet, causing the music to pause while waiting. TCP on the other hand provides a reliable stream of data. This protocol is a well known protocol that most of Internets traffic uses. When sending an e-mail you want all of the data to get through since half of the mail is of no use and you also want the data to arrive in order.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

When using TCP the first thing to bear in mind is that the TCP timeout is set to three times the round trip time (RTT). The reliable channel will however try to re-send the lost packet more times than that. This means that in the beginning the TCP transfer will not suffer from the backoff effect caused by the congestion control mechanism. This is true if the packet gets through within three re-transmissions, after this the TCP timeout will trigger decreasing the ¨benefits of the reliable channel. When later on looking at the drawn charts the assumption is that if tests with a reliable channel whould have been conducted the plots would have looked the same in the beginning and as packetlosses starts to occur the plot using the un-reliable channel would starts decreasing while the reliable channel plot would have held its levels. As the resends starts to exceed three times the plot of the reliable channel would quickly drop to the same levels as the un-reliable channel. This would mean that the drop of the reliable channel would be much steeper then the un-reliable channel.

When using UDP over a un-reliable channel the effect is what you would expect, if the packet gets lost it stays lost. When using UDP over a reliable channel the effect would be that when we start loosing packets the baseband will try to re-send them until the Flush Timeout is reached, and the packet is definitly lost. If the packet get through before the Flush Timeout the baseband has waisted alot of bandwidth trying to re-send the packets. While this packet is retransmitted the following packets will be qued in the senders out-buffer. The effect of this behaviour would be that we would see a decreace in throughput caused by decreased send rate at the sender. The receiver will not experience any lost packets as log as the packet gets delivered before the Flush Timeout, hovewer the receiver will see increased jitter. The packet can be received between the first and Flush Timout –1 transmissions. When sending DH5 datapackets each packet takes five timeslots to send and one timeslot to acknowledge all together six time slots. Each timeslot i fixed to 625 microseconds. If Flush Timeout is set to ten RTTs it will take 625*6*10 = 37500 microsec. = 37.5 milliseconds before the packet is considered as lost. In computer time, 37.5 ms. is quite much.

3.3 IEEE 802.11b, Wireless LAN

This standard, as indicated, defines the physical and medium access layer adapted to the requirements of a wireless network, but it offers the same ‘higher layer’ interface as the 802.3 specifications. The wireless 802.11 fits seamlessly into a wired 802.3 network. The only differences that a user might experience are lower bandwidth and sometimes longer access time. The mobile devices, called stations, connect to the wired Ethernet via the so called access points. The access points works as a gateway between the wired and the wireless network providing access to many WiFi stations at a time.

In order to be worldwide operable the standard uses the 2.4 GHz ISM band. This band is free of license and is also used by the competing radio technique Bluetooth. 802.11b, a.k.a Wi-Fi™, specifies a mandatory bandwidth of 1 Mbits/s and the optional 2, 5.5 and 11 Mbits/s. It uses direct sequence spread spectrum (DSSS). The sender applies a chipping sequence consisting of smaller impulses to every sent user bit. This means that a bit will be split into several smaller pieces. IEEE 802.11 uses the so-called Barker Code, which is robust against interference and is insensitive to multipath propagation.

The medium access protocol supports Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) with ACK. The easiest explanation of the operation would be that you “listen before talking”. This technique does not eliminate collisions caused by simultaneous transmission by multiple devices, but it eliminates sufficiently many of them. The difference between a wired and a wireless network is that when you use radio you ca not detect a collision while sending, thus the sender will not know that its transmission failed until the timeout of the ACKnowledgement.

IEEE 802.11 devices offer power management support. Standard LAN protocols assume that the receiver is always ready to receive data, but for lightly loaded networks the receiver is

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

mostly idle. According to the specification, the Wireless LAN card's power consumption in receive mode is 185 mA, but when the card is in sleep mode the power consumption drops to 9 mA. This power save is gained by turning off the transceiver whenever it is not needed. But since you cannot know when it should receive packets it has to wake up at regular intervals. These intervals have to be synchronized among all the stations and access points so that they are awake at the same time. When a station is in sleep mode the applications might still want to send data. It is then the job of the card to buffer this out-going data transparently to the application and send it as soon as the card goes to awake mode. The access points use the same technique, which buffers the data going to the stations while they are in sleep mode. So to summarize, what you gain on lower power consumption you will lose in bandwidth.

3.3.1 Wi-Fi

Wi-Fi™ is the abbreviation for Wireless Fidelity, which describes the IEEE 802.11b standard. The non-profit organization WECA (Wireless Ethernet Compatibility Alliance) was formed in 1999. The goal is to certify interoperability of Wi-Fi™ products and hence promote Wi-Fi™ (IEEE 802.11b) as the global wireless LAN standard.

In order to do this the WECA has defined a test suite that defines test that products have to pass in order to prove they are compatible with the standard and other certified products. At the moment (2001/04/18) more than 80 companies have joined the WECA. A press release [Weca2000] from 2001/05/07 announces that more than 112 products from 43 different companies have passed the test and thus received a Wi-Fi™ certificate.

3.3.2 PCMCIA-card from Lucent

The test cards used are Lucent ORINOCO PC-Card Silver with WEP encryption. These cards have been Wi-Fi™ certified and follow the European ETC standard.

3.3.2.1 Technical specification

The test cards used were the Wi-Fi™ certified Lucent ORINOCO PC-Card Silver [ORiNOCO2000]. These cards use Wired Equivalent Protection (WEP) with a 64-bit key. There's a gold version of these cards with a 128-bit key, using RC4 encryption in available in the USA.

As with Bluetooth the Wi-Fi™ operates in the license free 2.4 GHz ISM band. The modulation technique is Direct Sequence Spread Spectrum (DSSS) with the 11-chip Baker Code. These cards supports 1 and 2 Mbits/s transfer rates, and optionally 5.5 and 11 Mbits/s.

The maximum output power level is 100 mW (EIRP) in Europe and 1000 mW in the USA. These levels are regulated in the ETS 300-328 and the FCC 15.247 documents [802.11, section 15.4.7.1, p.234].

3.3.2.2 Availability and price

The WaveLAN cards from Lucent ORiNOCO are available from 1600 SEK + tax. There are over 15 other companies selling their Wi-Fi products in Sweden with prizes starting at 1300 SEK + tax.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

3.4 Reference data rates of the two techniques

3.4.1 Bluetooth

Packet Type	Max. Payload (bytes)	FEC	Max Asymmetric Data Rate forward
DM1	17	2/3	108.8
DH1	27	None	172.8
DM3	121	2/3	387.2
DH3	183	None	585.6
DM5	224	2/3	477.8
DH5	339	None	723.2

Bluetooth specification provides a variety of different packets for carrying data. For ACL links DM and DH packets are of interest. DM stands for Data Medium rate and DH stands for Data High rate. DH packets reaches higher speeds because it uses less error correction in the packet, this makes it possible to carry more data in each packet and thereby increasing the throughput. Data packets can span one, three or five time slots. Naturally a 3-slot packet carries more data than a 1-slot packet and a 5-slot packet carries more than a 3-slot packet.

This work only covers the behavior when sending data and thus using Bluetooth asynchronously and connectionless. This means that the maximum expected transfer speed couldn't exceed 723.2 Kbits/s.

3.4.2 Wireless LAN

Wireless LAN - 802.11b - provides the user with data rates from 1, 2, 5.5 and 11 Mbits/s. According to ORiNOCO the effective range at 11 Mbits/s is 160 meters and at 1 Mbits/s 550 meters in open space. In closed offices the effective range is 25 meters at 11 Mbits and 50 meters at 1 Mbits/s [Orinoco2000]. Cards supporting multiple data rates can perform dynamic data rate switching with the objective of improving performance. The specification doesn't provide any explanation of how this is supposed to be done and when [802.11, section 9.6, p. 111]. This behavior can be seen in the tests.

Powersave mode might have a negative affect on the throughput. When a mobile station (STA) goes into powersave mode it informs the Access Point (AP). The AP should then start to buffer packets to the STA and packets to the STA at designated times that informs the STA that it has information in the AP. The STA has to wakeup periodically and listen for these beacons that tells it that it has data stored at the AP.

4 Experimental Variables

The objective is to study how interference between Bluetooth and WiFi will affect the throughput. Varying the distance between Bluetooth and WiFi equipments varies the level of interference.

As briefly mentioned earlier there are numerous aspects that have to be taken into account when switching from a wired to a wireless network infrastructure. Some of these aspects concern the radio interference from other equipment. Among those lies the problem with interference from other competing radio techniques or radio noise from machines emitting radio signal in the 2.4 GHz frequency band. Ortivus will use Bluetooth for transferring data between the measurement unit and the PWS. One possible scenario would be if the hospital or another company in other environments where using products from both Agilent and Ortivus in the same room or that the family has one or probably more wireless equipment. Ortivus and Agilent products do almost the same job but they use different kinds of radio transceivers, Bluetooth and Wireless LAN.

There are two possible ways to go when you try to gain knowledge of how interference between WiFi and Bluetooth influences the behavior of wireless networks. The first is simulations, previously performed by Mobilian Corporation [Mobilian2000 & Mobilian2001]. Measurements, has not been done to the same extent, possible due that it is harder to get confidential results, tests takes longer time to carry out and the results are more equivocal than those given by simulations.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

My work is based on measurements performed in the vicinity of the Polaksbacken campus, Uppsala, Sweden in the spring of 2001. The tests were performed both indoor and outdoor.

The test scenarios consist of UDP and TCP throughput measurements at different signal strengths with or without interference from the other radio sources. In the tests both FTP transfer with the program CuteFTP from Globalscape and the bandwidth measuring tool Iperf where used. Iperf is described in more detail in the appendix A.

The goal was to study how the throughput decreases with increased distance and interference from other radio sources.

4.1 Measurement metrics

4.1.1 dBm

The signal strength is measured in a unit called dBm, where the 'm' stands for milli-watts of power. dBm is simply referenced to as milli-watts where 1 mW = 0 dBm. For instance if we have +7 dBm, (following the logarithmic seal rules) and dividing +7 by 10 we get 0.7 and the anti-log of that is 5.0118. This means that +7 dBm is another way of saying a gain of 5 milli-watts. The formula for calculating dBm is; dBm = 10 log (p/0.001) W. The reason for using dBms instead of milli-watts is that when you have systems with gains and losses it is much easier to calculate with dBms.

4.1.2 Signal-to-Noise Ratio (SNR)

Signal-to-Noise Ratio is the ratio of the signal power to the noise power. This ratio is expressed in decibels (dB). For instance if a link has a SNR of 50 dB, this means that the signal power at the output is 50 dB above the noise power.

For the Lucent ORiNOCO Wi-Fi™ card, SNR values over 20 dB is ranked as from good to excellent link quality. Values could go up to as much as 85 dB at very close distances. At 2 meters the SNR are typically between 50 and 55 dB. SNR levels between 20 and 10 dB is considered as marginal link quality and below 10 dB means that the signal strength is "Out-of-range". When the SNR drops below 5 dBm the link is terminated.

The Link manager program, which comes with the Lucent ORiNOCO Wi-Fi™ card, provides information about the quality of the link. This program makes it possible to monitor SNR value, signal strength and noise levels between the test partners. The program sends four packets per second between the test partners. For every packet the SNR, signal strength and the noise level are calculated. This program could also be used to retrieve a list of all basestation that are in reach at the moment with information about the link quality.

The Bluetooth suite doesn't provide any means of extracting the SNR, signal strength or noise level. Digianswer has announced that the next driver will support the 1.1 specification of the Bluetooth standard and thus support monitoring of the signal strength and noise level. It is not stated when a new GUI will be released that provides this information. At the moment (2001/06/04) the driver implementation supporting the 1.1 standard is in beta-stage and therefore quite unstable. The specification gives the minimum performance parameters for the RF system. According to the specification Bluetooth should operate with a maximum Bit Error Rate of 0.1%. This gives receiver sensitivity of –70 dBm. But most devices exceed this value with 10 dBm or more.

4.2 Problem with signal propagation

The received signal strength is proportional to $1/d^2$ with d being the distance between the sender and the receiver. The explanation for this is that the sender emits a signal with certain energy. This signal travels from the sender in a spherical shape, if nothing blocks its path the signal strength decreases as a function of $S = 4\pi d^2$. This formula is valid for vacuum, and it is good enough for air.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

When dealing with short distances like LANs the attenuation is therefore not so much of a problem. There are other problems.

4.2.1 Problems

The wavelength of a transmission in the 2.4 GHz band is approximately 0.125 meters. (Wave length = speed of light /frequency). When a wave hits an object that is smaller or in the size of the wavelength the waves can be **scattered** into several weaker signal bouncing away in different directions. Another similar problem is **diffraction** of waves. This means that the waves can be deflected as they hit an edge and propagate in a different direction.

The worse problem is **multipath propagation**, when waves that have bounced arrives out of phase with the direct wave and thus partly canceling the signal. The bouncing causes **delay spread**: the original signal is spread due to different delays of parts of the signal. This means that shorter impulses will be smeared out into a broader impulse, or rather several weaker impulses. These weak impulses have taken different paths and have different attenuation, thus the received pulses have different power.

This effect is called **intersymbol interference** and limits the bandwidth of a channel with multipath propagation. The result of scattering, diffraction and multipath propagation generate varying signal strength depending on the location of the receiver.

The interior walls at Polacksbacken buildings are made of thick concrete. This means severe loss of signal strength. However some energy bounces on the walls, which makes it possible to reach around corners. But this wall bouncing causes the signal level to fluctuate a lot. Our tests show that moving the receiver just a couple of centimeters sometimes results in a 5-10 dBm drop in the signal strength. If movements are done at low signal strengths the link could easily be dropped.

Another problem I have experienced is that people moving in the area between the basestation and the measuring computer affect the signal strength. The body shields or absorbs the signals and thus affects the link quality. It was hard to quantify how much the movement of people affected the signal strength and our measurement were collected when there were no people around.

When close to the sender the relative signal strength fluctuates more compared to when you move the receiver near the middle of the covered area.

4.3 Distance vs. lost packets

The signal strength falls off rapidly as the distance between the sender and receiver increases, roughly as a function of $1/r^2$ in air, i.e. the quality falls quickly. The link will become very unreliable if we reside in these areas and the throughput will drop. When the ratio between the signal strength and the noise gets below a certain value the receiver will be unable to tell the difference between the signal and the noise. When this happens the receiver can no longer reconstruct the packets, which will be dismissed. Under these circumstances the higher layer protocols can either improve or aggravate the throughput.

4.4 Influence of movement

The situation gets worse when the receiver and/or sender move. When the devices move, the channel characteristics change over time as well as the received signal strength at the receiver. These quick changes in received signal power are called **short-term fading**. To compensate for this variance the receiver tries to adapt to the new conditions, but if the changes are to quick the receiver might not be able to adapt, thus causing errors in the transmission.

There is another effect, called **long-term fading**, which could be described as the received signal strength average over a time period. Changing distances or introducing obstacles might cause this effect. To compensate for this the sender can adjust the sending power so that the receiver gets a signal.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

As soon as you start moving the device, the signal strength drops by several dBm. Measurements show that the drop is approximately between 3-6 dBm. However the drop seems to be lower at closer distances. For the measurement it takes time before the signal strength stabilizes. My guess is that is due to a combination of both multipath propagation that causes short term fading and that the measuring unit in the Wireless LAN card from ORiNOCO card might display values based on an average over some time period that might not be valid anymore when on the move.

4.5 Interference from competing radio techniques

Radio transmission in the ISM band is exposed to interference from other radio techniques using the 2.4 GHz band since it is globally available. But there are other types of equipment that emit in this band. The most common is microwave ovens. A modern microwave oven normally consumes some 1000 Watts. If just a small amount of this transmission escapes out of the oven it could easily drown a Wi-Fi or Bluetooth transceiver.

Cordless phones, such as the DECT standard, uses the 1880 – 1900 MHz band and should not pose any problem with interference to devices using the 2.4 GHz band.

Radar signals can range from 500 MHz to 9 GHz depending on which purpose it has. The probability of interference from radar when using ISM products is in my opinion quite low.

4.6 Method

Our objective is to measure the Bluetooth and Wireless LAN links under interference. Since we cannot observe the signal strength of the Bluetooth link the first step will be to measure the Wireless LAN connection. We first measure a reference throughput. In our figure the y-axis shows the throughput and the x-axis shows the measured signal strength at locations with increasing distance from the access point. Lines will be plotted for both TCP and UDP throughput since the assumption is that the behavior between them differs. UDP will provide us with information about the losses. TCP is interesting because most standard applications use the TCP protocol that has error recovery and is more complex than UDP. The same tests will be done with Bluetooth but since we cannot see the signal strength of the Bluetooth link corresponding geographical positions has been used. These positions are marked with letters ranging from A to G on the map below in figure. At these positions the signal strength of the Wireless LAN was measured. Then the Wireless LAN basestations and pc-cards are shutdown and the Bluetooth master is placed at the location of the Wireless LAN basestation outside room 1421 and the Bluetooth slave is placed at each of the marked locations.

When doing the reference for the Wireless LAN a test with powersave mode active was performed, to be able to see how much this mode affects the throughput.

When we have the above described reference lines the same tests will be done again, but this time interference from the other technique will be added between the measuring units in order to see how the link behaves and how much the throughput drops. The interference will be added on a “most likely to happen” basis and a worse case scenario. The scenarios differs somewhat between the two techniques the reason for this is the difference between the two techniques and the usage scenarios we anticipate.

When testing the Bluetooth performance the networking profile will be used since this profile gives the highest throughput and is the most likely to be used in a telemedicine application.

4.7 Test environment

4.7.1 Computers

The tests where conducted on three different laptop machines:

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

- Toshiba Tecra 8000: Windows 98, Intel Pentium II 300 MHz, 128 Mb RAM
- IBM 600: Windows 2000, Intel Pentium II 233 MHz, 64 Mb RAM
- AST Ascentia P series: Windows 98 and Linux, Intel Pentium 133 MHz, 24 Mb RAM

These machines will be referred to as Toshiba, IBM and AST.

4.7.2 Test locations

The measurements were carried out at the old regimental buildings at Polacksbacken in Uppsala, Sweden. These houses have thick walls that effectively block the radio signals from spreading horizontally, and thin floors that allow the signals to relatively easy propagate vertically. The test were indoors under the coverage of a Wireless LAN with eleven base stations. The base stations are positioned such that they maximize the coverage.

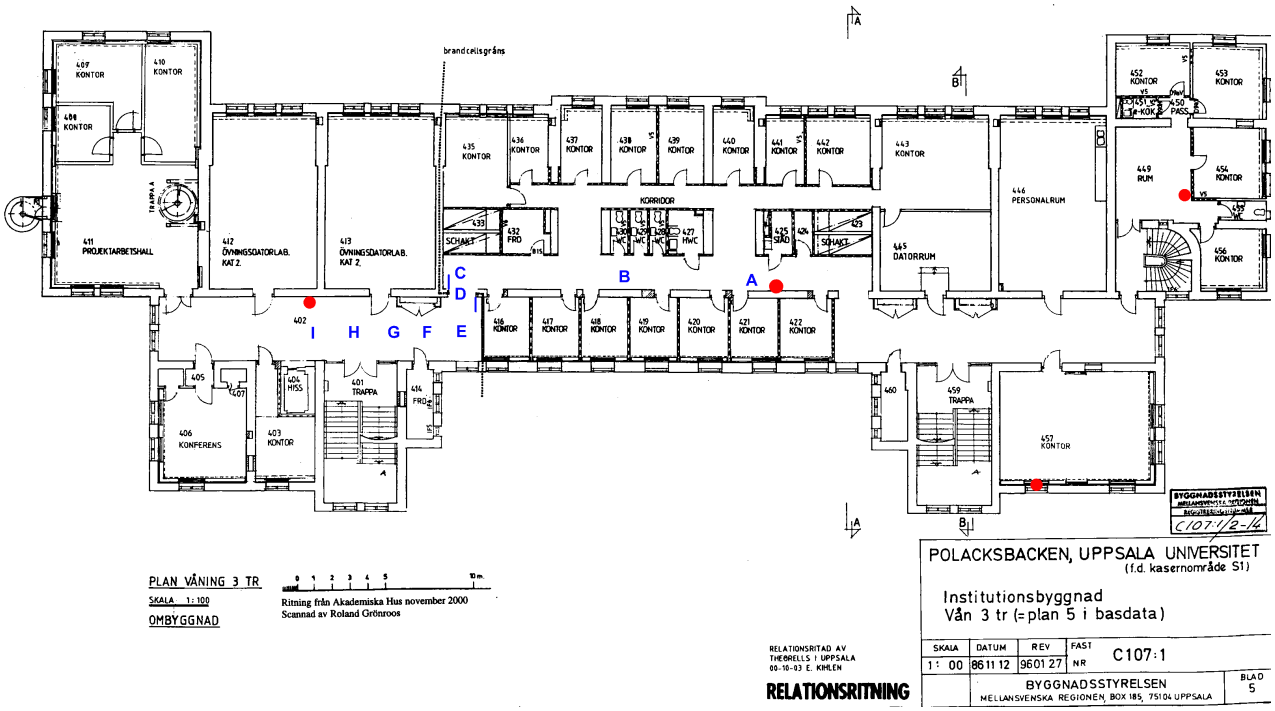

Figure: Polacksbacken campus building 1 floor 4

Figure above describes a map of floor 4 in house 1. The circles are wireless LAN basestations from Lucent with a small range antenna. There are also basestations in the basement, on floor one, two and five. While doing the measurements the basestations that were not used where shutdown for the duration of the tests.

The letters on the floor, ranging from A to I, are the spots where Bluetooth measurements were conducted. The Bluetooth master is placed at the same position as the Wireless LAN basestation located outside room 1421.

Every point in the following graphs is the average of two to seven independent test runs where each test run is the result of FTP'ing a 20 Mb file or running Iperf for 30 seconds. When testing Bluetooth the FTP'ed files were smaller. When the two test devices was close to each other there was no need for more than two independent tests, as the distance between the two devices increased the need for more tests increased since the variance increased.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

5 Interpretation of results

Note that in the following charts the distance between the points -40, -50, -60 and -70 are the same as the distance between the points -75, -80, -85, -90, -95 and -100. To compensate for this picture the distances between –75 and –100 as half the distance. This will make the curves drop much steeper than it actually looks like.

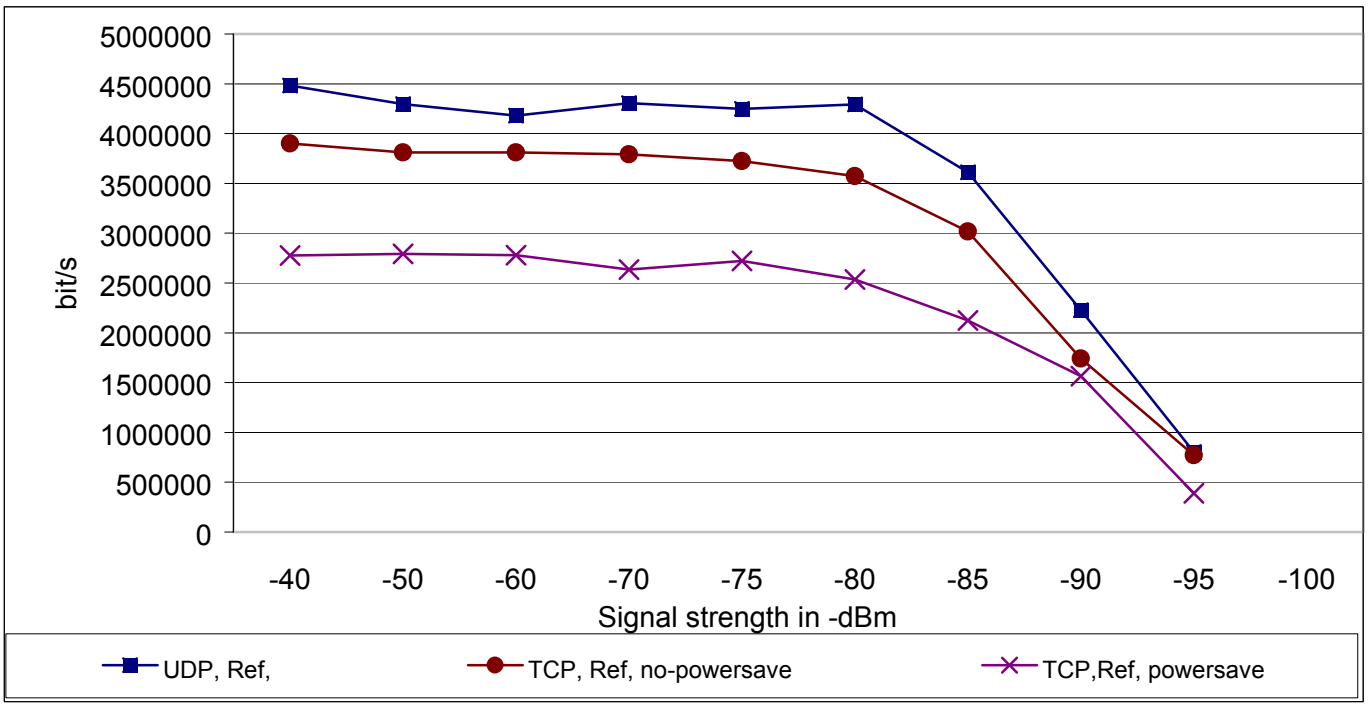

WiFi: UDP, TCP and TCP with power save

In chart the plotted lines shows three reference measurements with the Wi-Fi card. The squared line shows a UDP transfer and the circled line shows the same measurement using TCP. The X'ed line shows the TCP transfer with the power save mode active. The y-axle shows the throughput in bits/second and the x-axle shows the measured signal strength in dBm. Between –40 and –70 dBm the measurements where done every tenth dBm and between –75 and –100 every fifth dBm. (This makes the lines drop smother than they actually are.) It is also possible to measure the throughput at lower dBms than –40 but it is not meaningful since the focus is not to see the behavior when the devices are near each other and the link is good. The interest is to see how the throughput decreases.

At –95 dBm it is hard to get reliable results as the throughput varies. When I did the measurements it was possible to have access at –100 dBm but since the link was so unstable it was impossible to do tests at this value because the link dropped quite often during on test run, this is why results from this dBm value is not included.

The line on top shows the UDP performance and the line in the middle shows the curve using the TCP protocol. What we can see is that UDP starts with a higher throughput than TCP and it keeps this advantage until the signal strength drops to –95 dBm. The reason for this has to do with both the fact that UDP has less header as thus uses the packet more efficiently and the fact that UDP performance doesn't fall like TCP when packet losses occur.

Comparing TCP without power save mode and TCP with power save mode that the performance gets affected if power safe mode is active. The assumption is that the power save mode line would drop steeper because of the TCP congestion control that decrease the send rate when the connection gets bad. Packets that are stored in the AP will make the TCP think that there is congestion and thus decrease the send rate.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

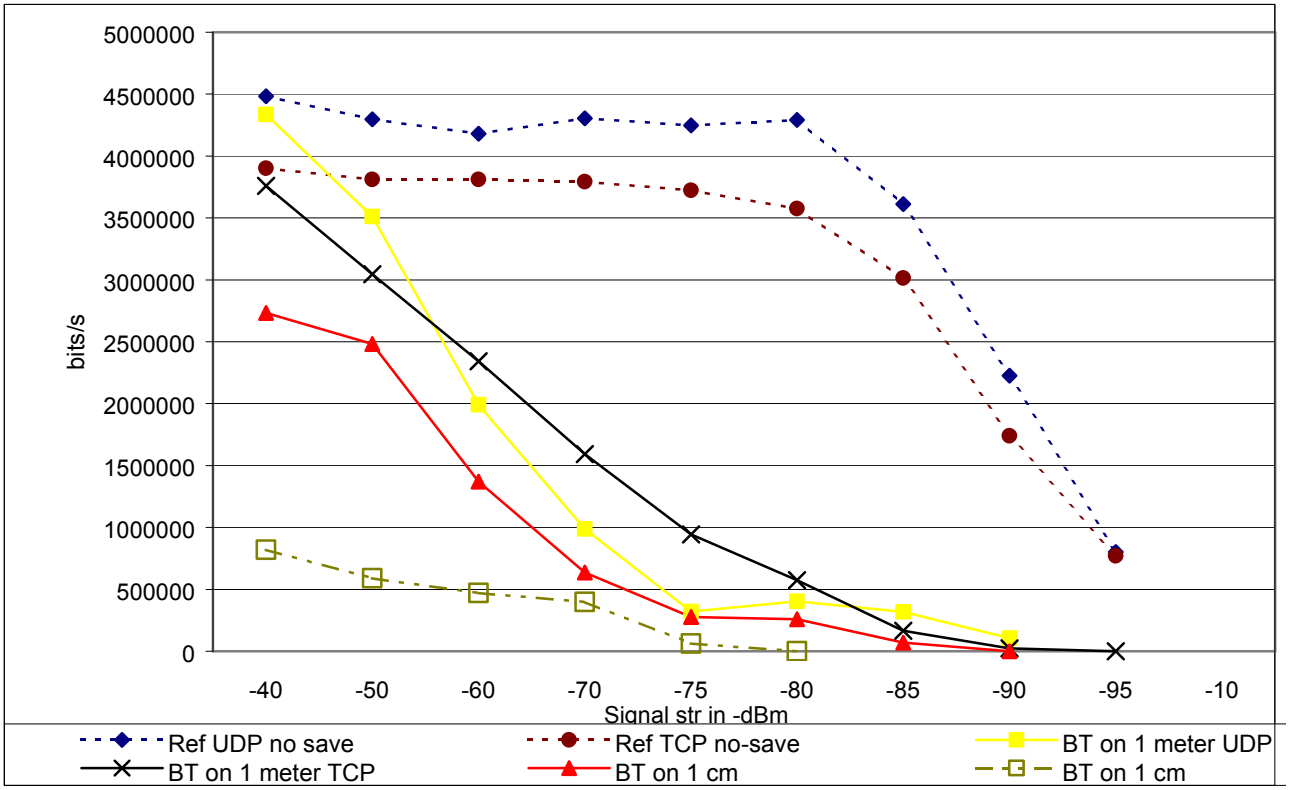

Chart: Added Bluetooth interference

5.1 Bluetooth interference

In this chart the TCP power save line is removed and the previous two are drawn with dashed lines. The four new lines show Wireless LAN exposed to interference from Bluetooth devices.

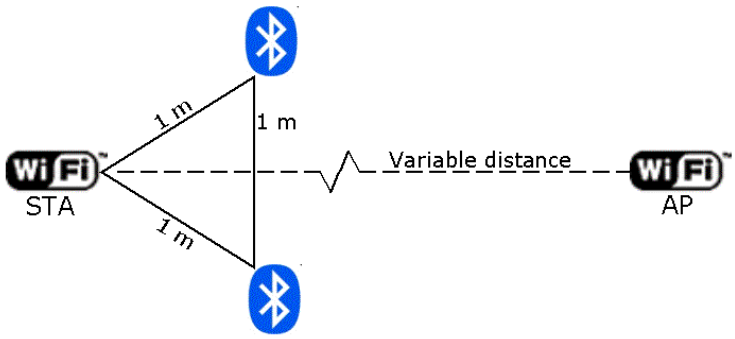

Picture: Bluetooth interfering with WiFi

The two Bluetooth devices are set up between the two Wireless LAN devices one meter from the mobile Wireless LAN computer (STA). The computer and the two Bluetooth devices form a triangle with 60 degrees angles and one-meter sides, with the two Bluetooth devises in between the STA and the AP. The three devices are then moved away from the wireless LAN access point. At certain signal strengths throughput measurements are performed. The signal strength is the strength of the basestation measured at the moving Wireless LAN device.

The constellation is quite artificial but yet not taken from out in the blue. Bluetooth devices are anticipated to operate in close vicinity from each other whereas Wireless LAN has longer reach and are therefore expected to be used further apart. Bringing two Bluetooth devices close to one WiFi product is bound to happen for example when you sit in front of your portable laptop with Wireless LAN having your Bluetooth enabled PDA and mobile phone in your pocket.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

5.1.1 Bluetooth disturbing WiFi

In order to get the worse case scenario for the WiFi card the Bluetooth devices are sending UDP data from the slave to the master using the networking profile. This will make the Bluetooth devices use the media as much as possible and make sure there are no backing-off as they loose packets caused by the interference from the Wireless LAN devices. Making large file transfers or fast and long UDP transfers is not the most common usage-scenario for Bluetooth but since it is possible somebody might use it this way and thereby causing this scenario.

The yellow squared line shows the throughput of the UDP WiFi transfer and the X'ed black line shows the same test but with TCP. As can be seen the UDP line has a higher initial throughput than TCP, this is what can be anticipated when looking at the reference measurements. At –40 dBms the Wireless LAN does not get that much influenced by the interference caused by the Bluetooth devices. This probably has to do with the fact that Wireless LAN has a much higher output power than Bluetooth and thus floods the ether. As we come to the area between –60 and –80 dBms the TCP transfer shows higher throughput than UDP. This is quite strange and I have not figured out why this happens. The results are however re-producible. As we reach signal strengths over –85 dBms the UDP line crosses the TCP line and the UDP throughput is there after slightly better than with TCP.

The red line with triangles shows the Wireless LAN using the TCP protocol with the powersave mode active and the Bluetooth card in the slot under the Wireless LAN card but not active.

The green line with hollow squares illustrates what happens when you put the Wireless LAN card in one pcmcia slot and the Bluetooth card in the other slot. The initial and maximum transfer speed is less than the speeds measured at –95 dBm when there is no interference and only the Wireless LAN card present. A possible reason for this is that the Bluetooth card is located so near the Wireless LAN card that its transmission power floods the Wireless LAN. The Wireless LAN using TCP in this case backs of when it starts loosing packets. Bluetooth on the other hand seems to provide reliable connection at the physical layer and thus tries to re-send lost packets at the lower layer. This means that even if the Bluetooth device where using TCP the physical layer still utilizes the ether constantly in order to get the packets through. This setup is the worse case possible for the WiFi and the strong recommendation is that it should be avoided if possible.

The most common way of placing the PCMCIA slots today in laptops is to place them on top of each other. On the computer that was used in the tests the Wireless LAN card is placed just on top of the Bluetooth card. The cards are separated with only a couple of millimeters. This way of placing the PCMCIA slots in laptops will not be the standard way of placing them when these kinds of test results starts to show up. The most reasonable way of doing it would be to place the slots on opposite sides of the computer of inventing some other way of separating the antennas from each other.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

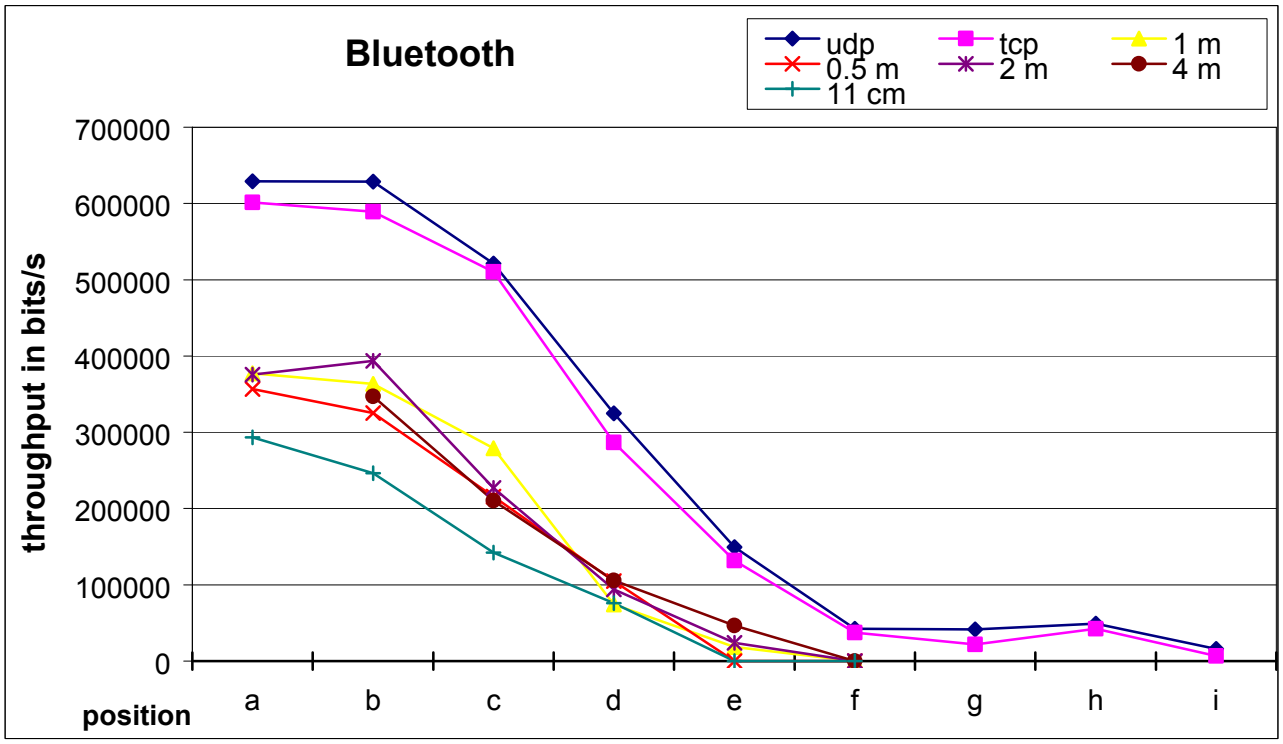

Chart: The Bluetooth measurements

5.1.2 WiFi interfering with Bluetooth

In this chart the focus is the Bluetooth behavior. As explained earlier in this document the test were conducted indoors at floor 4 at house 1 at Polacksbacken campus in Uppsala. The Y-axis shows throughput in bits/s and the *X-axis shows the Wireless LAN signal strength at the positions where we did the previous measurements*.

The two plots on top are the two reference lines were the transmissions were not exposed to interference from any other radio source. The four lines below show what happens when a computer with a Wireless LAN card starts interfering between the two Bluetooth devices. To get a maximum of interference the Wireless LAN card does a ping flood-broadcast with the maximum Ethernet packet size. This scenario will make the transmission time as large as possible and thus maximizing the interference.

The wireless LAN transmission covers 22 of the 80 one-MHz frequencies in the ISM band. The interference the Bluetooth device is exposed to is limited to those hops when it lands in the area where the Wireless LAN device is sending. The Wireless LAN network uses 22/80, which is 27 percent of the available band that Bluetooth use.

The table to the right shows both the TCP and UDP throughput at 1 meter with added interference divided with the reference throughput values. The table shows that we loose at least 100-63=27% of the capacity when we add interference from Wireless LAN this is what we can expect from the calculation above. As the Bluetooth link gets weaker the interference from the Wireless LAN influences the connection more and finally the connection will be lost both due to the distance and the interference.

Position	TCP	UDP
A – 49	63%	60%
B – 52	62%	58%
C – 69	55%	54%
D – 78	26%	23%
E – 80	14%	12%
F – 82	0%	0%

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

6 Conclusion

Both Wi-Fi and Bluetooth use the ISM band. In the context of telemedicine the two techniques complement each other. Bluetooth offers mainly a cable replacement even though the most eager Bluetooth promoters want it to perform more complex tasks. Wireless LAN provides networking capabilities like a wired Ethernet. My prediction is that both techniques will have strong market positions in the future. However the foreseen price of the Bluetooth equipment of about five dollars a piece, and not like today with 150 dollars for a pc-card, will only come true if the sales reach the predicted quantities of several million devices produced and sold every year, which will bring down the production costs, thus reducing the prices. If the same growth estimates where applied to Wireless LAN the prices for these cards would probably drop to a similarly low level. If Wireless LAN had the same price as Bluetooth price my choice would easily fall on the Wireless LAN since it is more capable than Bluetooth at the moment; I am referring to the better networking abilities and the greater transfer speeds.

Bluetooth and Wireless LAN are at the moment the two major techniques in the ISM band but since this band is license free there is the possibility that new techniques and products will emerge in the future. This is not something that we cannot say so much more about at the moment. The interference problem between Bluetooth and Wireless LAN is discussed in IEEE 802.15 standard. The IEEE 802.15 develops Personal Area Network consensus standards for short distance wireless networks; a.k.a WPANs™. The task group 1, TG1, focus on the co-existence of Bluetooth and Wireless LAN. More information at: http://www.ieee802.org/15/pub/TG1.html

Based on the measurements, we conclude that both Wireless LAN and Bluetooth gets affected of interference from radio transmissions in the ISM frequency band, the Bluetooth devices have longer range than the ten meters the specification indicates, and when you add interference the throughput decreases rapidly. However if the demand on the throughput only were 100Kbits/s both techniques would be sufficient.

My conclusion is that Bluetooth is not mature enough for telemedicine implementations; the primary reason for this is that the interference from Wireless LAN or other radio techniques has not been studied enough. I am also uncertain about the use of Wireless LAN. But the advantage of wireless LAN is that the initial throughput is higher which makes it take longer time to reach a level where the available throughput is lower than needed. The other reason for this is that the software based on the current Bluetooth specification 1.0b doesn't support monitoring of the signal strength, noise level and SNR. Without this feature it is impossible to make programs monitor the connection quality and set alarm levels when the link gets bad or when the connection is lost.

Wireless LAN suffers from poor safety caused by bad design and more often faulty implementations of the WEP protocol that makes the security in the Wireless LAN network less than satisfactory. During the last half year there has been a flood of research reports showing the bad design of the WEP protocol.

Copyright © 2001, Magnus Berggren

7 References

[802.11 1999] ANSI/IEEE Std. 802.11, 1999 Edition (20/8/1999) also ISO/IEC 8802-11:1999(E): *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. ISBN 0-7381-1658-0, SH94740. (25/6/2001.) http://standards.ieee.org/getieee8o2/802.11.html

[801.11b 1999] ANSI/IEEE Std. 802.11b, 1999 Edition (16/9/1999): *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band*. ISBN 0-7381-1812-5, SS94788. (25/6/2001.) http://standards.ieee.org/getieee8o2/802.11.html

[Agilent2000] Agilent, Oct. 5 2000: Press release, Agilent Technologies Announces Availability of Wireless Networking Capability for Portable Patient Monitoring. (15/4/2001.) http://www.healthcare.agilent.com/press_releases/PRHS2920030.html

[Baquet1997] Baquet, C.R, 1997: An Overview of Telemedicine In: *Journal of the Association for Academic Minority Physicians*, Vol. 8(1997):1, page 2-10, Baltimore, Md. / AAMP. (4/9/2000.)

[Bray2001] Bray, Jennifer & Sturman, Charles, 2001: *Connect Without Cables*. Prentice-Hall, Inc. NJ 07458. ISBN 0-13-089840-6.

[Ferrer-Roca1998] Ferrer-Roca, Olga & Sosa-Iudicissa, Marcelo, 1998: *Handbook of Telemedicine*. IOS Press, Amsterdam.

[Istepanian1998] Istepanian, Robert H., 1998: Design of mobile telemedicine systems using GSM and IS-54 cellular telephone standards, In: *Journal of Telemedicine and Telecare*, Vol. 4(1998), Supplement 1, page 80-82.

[Karp2000] Karp, Warren B, 2000: Use of Telemedicine for Children With Special Health Care Needs. In: *PEDIATRICS*, Vol. 105 No. 4 April 2000. ISSN 0031-4005. (20/9/2000.)

[Löf2000] Löf, Jonas, 2000: Erfarenheter från ett telemedicinskt pilotprojekt: Prehospital Ambulanssjukvård. http://www.tmc.hs.se/pub/Rapporter/PreHospAmbulans.pdf (19/9/2000.)

[HP 1/1998] Hewlett-Packard, 1/1998: *Communication Standards in Clinical Setting: An Introduction*. Publication No. 5966-3748E

[HP 4/1998] Hewlett-Packard, 4/1998: *Communication Standards in Clinical Setting: A Closer look at the Details*. Publication No. 5967-5483E

[HP 5/1998] Hewlett-Packard, 5/1998: *Communication Standards in Clinical Setting: The Hewlett-Packard Solution*. Publication No. 5967-5827E

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

[Holm1999] Holm Sjögren, Lotta, Sandberg Carolina & Schweiler Åsa, 1999: *Telemedicine Applications in Sweden 1998*. ISSN 1403-3356. Sprinternet 6. (5/2/2001.) http://www.spri.se/frlag/elektron/Sprint6.pdf

[Main2000] Main, Michael L. et al, 2000: Real-Time Transmission of Full-Motion Echocardiography over a High-Speed Data Network: Impact of Data Rate and Network Quality of Service, In: *Journal of the American Society of Echocardiography*, Vol. 13(2000):8, page 764-770

[Medline] Medline, NCBI. http://igm.nlm.nih.gov/cgi-bin/doler?account=++&password=++&datafile=medline or http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?CMD=search&DB=PubMed

[Mobilian2000] Mobilian Corporation, September 2000: *Wi-Fi™ (802.11b) and Bluetooth™ Simultaneous Operation: Characterizing the Problem.* (21/3/2001.) http://www.mobilian.com/documents/Characterizing_the_Problem.pdf

[Mobilian2001] Mobilian Corporation, January 2001: *Wi-Fi™ (802.11b) and Bluetooth™: An Examination of Coexistence Approaches*. (21/3/2001.) http://www.mobilian.com/documents/Coexistence_of_802.11b_and_Bluetooth1.pdf

[Mobimed] Mobimed – Decision making Support System in Pre-Hospital Emergency Care. Ortivus, 1999. (14/11/2001.) http://www.ortivus.se/mobimed.asp?mi=prod&submi=mobimed

[Orinoco2000] ORiNOCO, 19/5/2000: ORiNOCO PC Card Specification. (1/6/2001.) ftp://ftp.orinocowireless.com/pub/docs/ORINOCO/BROCHURES/O_PC.pdf

[Ortivus1999] Ortivus, 1999: MIDA monitoring: an overview.

[Phillips1999] Phillips, Michele, 1999: Telemedicine in Neonatal Intensive Care Unit, In: Pediatric Nursing, Vol. 25(1999):2, page 185-186, 189

[Salvemini1999] Salvemini, Anthony V. Ph. D., 1999: Challenges for User-Interface Designers of Telemedicine Systems, In: *Telemedicine Journal*, Vol. 5(1999):2, page 163-168.

[Schiller1999] Schiller, Joshen, 1999: *Mobile Communications*.

[BT 1.0draft 1999] Specification of the Bluetooth System, version 1.0 draft, 5/7/1999. (15/10/2000.)

[BT 1.0B Core 1999] Specification of the Bluetooth System, version 1.0B Core, 1/12/1999. (9/10/2000.)

[BT 1.0B Prof. 1999] Specification of the Bluetooth System, version 1.0B Profile, 1/12/1999. (9/10/2000.)

Copyright © 2001, Magnus Berggren

[BT 1.1 Core 2001] Specification of the Bluetooth System, version 1.1 Core, 11/2/2001. (2/5/2001.) http://www.bluetooth.com/developer/specification/Bluetooth_11_Specifications_Book.pdf

[Tanenbaum1996] Tanenbaum, Andrew S, 1996: *Computer Networks. Third Edition*. Prentice-Hall, Inc. NJ 07458. ISBN 0-13-394248-1.

[Telemedicine1999] Telemedicine: An Overview, In: *Health Devices*, Vol. 28(1999):3, page 88-103

[Vasco2000] Vasco Advisers, 2000: *Utvärdering av SABH*. (19/6/2000)

[Weca2000] WECA, 2000: Press release. Wi-Fi™ Passes 100 Product Mark – Far exceeding Industry Expectations And Competing Wireless LAN Standards. San Jose, CA, May 7 2001. (15/5/2001.) http://www.wirelessethernet.org/pressreleases/20010507_allcertifications.asp

[WHO1997] WHO, 1997: Press release. WHO Director-General Highlights Potential of Telemedicine. WHO/65, 16/9/1997, http://www.who.int/archives/inf-pr-1997/en/pr97-65.html (15/9/2000.)

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

8 Appendix

8.1 A, Sniffer programs

A part of this job has been to look at different packet capturing programs. There are numerous to choose from on the market, ranging from free to prizes exceeding $1000. Educational institutions always have a chance of getting free licenses.

8.1.1 Commview 2.4

CommView is a network-monitoring program capable of both capturing and analyzing packets on any Ethernet network. CommView lets you see a list of network connections, IP statistics as well as analyzing of individual packets. Packets can be accessed down to the lowest layer and stored to a log file for later analysis. During capturing, a flexible filter system could be used to drop unwanted packets.

The application is designed up to medium-sized networks and runs on all Windows systems. It requires an Ethernet network card supporting the NDIS 3.0 driver standard.

The evaluation version of the program is limited to a 30 days trial, which only displays every second packet. There are three different registered versions the Personal version costs $49, Professional $79 and Professional MultiOS $119. The personal version can only see packets going to or coming from the computer from which it is installed. The Professional version cannot be used on Windows NT/2000. The Professional MultiOS has none of the previous limitations. TamoSoft where kind and gave me a MultiOS license, which allowed me to use their program without limitations.

The program has an intuitive and compact interface. Just push the play button and the capture will start. It is easy to change which network device you want to listen to using a drop-down combo box containing all active devices. The statistics window gives you information about how many packets that are received per second and the average throughput of bytes per second these values are displayed both as a graph and as numbers. There is also a frame that tells you the total amount of packets captured.

8.1.2 NAI Sniffer

Network associates / Sniffer Technologies will release Sniffer Wireless in the near future. This program is built on the award-winning program Sniffer Pro and will offer full support to IEEE 802.11b, network monitoring, capturing, decoding and filtering. The program will have the ability to analyze frames from the Physical to the Application Layer. At the moment Sniffer Wireless integrates drivers from Symbol Technologies and Cisco Wireless LAN cards, but rumors tells that there will be more cards supported in the future.

8.1.3 Analyzer, WinDump and WinPcap

These programs are made by the netgroup at Politecnico di Torino, Italy, which can be found at URL: http://netgroup-serv.polito.it/analyzer/. They are released under a BSD-styled license. The documentation is often in Italian, and thereby somewhat hard to understand.

WinPcap is an architecture for packet capture and network analysis under Windows, based on the Berkeley Packet Filter and libpcap for UNIX. It includes a kernel level packet filter driver, a low-level dynamic link library (packet.dll), and a high level and system independent library based on libpcap, version 0.4a6. Both of the other programs uses WinPcap to capture and handle network traffic.

Analyzer had problems with opening the adapter I wanted to capture packets from. The only help I got was to reinstall the capture driver, which didn't solve the problem.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

WinDump is a fully compatible port of the UNIX program tcpdump, version 3.4a6. The program runs under DOS with the same commands as tcpdump, to analyze the results an external plot program could be used. When sending pings to some host WinDump only shows the ping-replies not the requests when you are listening on the same machine as you are sending on.

8.1.4 TTCP, PCA_TTCP and NTTCP

Both of these Test TCP programs are based on the version developed in 1984 for the BSD operating system. Test TCP (TTCP) is a command-line socket-based benchmarking tool for measuring TCP and UDP performance. To run the program, copy it to both of the computers in the test. One of the computers will be used as the receiver and the other will be used as a transmitter.

All of these programs are said to do the same job and being compatible with the original TTCP program. However I haven't been able to use one program on one computer and another program on another computer.

- PCA_TTCP can be found at URL: http://www.pcausa.com/Utilities/pcattcp.htm
- NTTCP can be found at URL: http://www.leo.org/~bartel/nttcp/

8.1.5 Iperf

Iperf is based on the original Test TCP program. Iperf has the ability to test a connection with both UDP and TCP data. According to the authors of the program Iperf should be compatible with TTCP in a way that you could use Ipref as a client and TTCP as a server or vice versa. I've had problems in getting this to work, but it could be the case that the implementation of TTCP that I've tried is not the same as the original TTCP version.

The best improvements when using UDP is that you can set the transfer rate at the sender, not as with TTCP where the sender always tries to send as fast as possible. The default bandwidth is 1 Mbits/s. Another option is that you could set either the amount of time the test should run or the number of packets to send. When measuring UDP Iperf not only tells which bandwidth that is available, it can also tell the jitter and lost packets. There is also the possibility to set the packet size.

When testing with TCP the best improvement is that you get information about the test while the test is running. Another good feature is that the Iperf server could be started as a service. If wanted reports can be given at different time intervals during the tests, giving a chance to see how the bandwidth changes when moving wireless devices during the test.

8.1.6 Agilent Analyzer, J1955A

The Agilent Analyzer is the most capable program I've tried in these tests. It has support for more than needed. However there is some minor thing that annoys me a lot. When using the packet generator the program is incapable of generating packets with sequence numbers in the packets this is of interest when sending data with UDP and you want to see which packets that got through. Another problem is that you cannot see the part of the packet you are interested in when you use the decode function. This function always displays the full packet. This means that even if you use a 21" screen the part of the packet you want to see is shown beneath want is shown on the screen.

According to Agilent whey are working on increasing the wireless support. At the moment only two Wireless LAN drivers are supported.

8.1.7 WS_ping ProPack

WS_ping ProPack has support for various interesting tasks. The version I tried where released in April 1998, version number 2.10. The most interesting with this program was its ability to use SNMP in an easy way. It has also support for, time synchronization, finger, TraceRoute, LDAP, ping, scan, et cetera.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

8.2 B, The client-server program

I wanted to study the characteristics of an application that behaves like the one that Ortivus is developing. The application should be able to send approximately 500 packets per second containing 64 bytes of data. At first I tried to use PCA_TTCP but due to the limitation of that program and the fact that there seems to be a bug in the Bluetooth driver or in the NDIS driver, I decided to develop a small client-server that could do what I want.

Since the Bluetooth PCMCIA-cards are sent without support for Linux, at this time only Win32 is supported, I was forced to make the program run under Windows. Since I have no experience coding in Windows I decided to use Java that is platform independent. But, occasionally there will come drivers that support Linux, and at that time the Java program will come handy.

The programs takes different flags and values as input some of them are optional and some are forced.

8.2.1 The client

Called: java UDPClient <-flags> ip-number|hostname

Flags:

-h	:This flag will print a helpmessage of how to use the program.
-n <i>	: The program will send this many packets and then exit. This flag is optional, and if not given a default value of 10 packets is sent.
-l <i>	: The packets will have a length of the value given. This flag is optional, and if not given a default length of 10 bytes will be sent.
-d	: This flags turns on debug printing.
-t <i>	: This flag tells the client to do a <i> millisecond wait between every package. This is use to be able the amount of packets that are sent per second. The default is no wait.
Ip/host	: This value could be a hostname as well and is not optional. The packets will be sent to the given hostname/ipnumber.

The program uses the UDP protocol to send packets to the server on port 4445, which at the moment is a limitation of the programs. If the debug flag is given the program will print the value of the packet to the screen. The packets contain a sequence number counting from 0, the time offset from the first packet and an empty string. These are separated by a ‘:’. After the termination of the program it will print the amount of milliseconds spent from the first packet to the last and the amount of packets sent.

At the moment the worse limitation of the program is that you cannot know when the server has terminated since the server do not send any acc’s. The simple solution I use is to send more packets from the client than the server except and let the server have some logic to take care of the problem.

8.2.2 The server

Called: java MyServer2 <-flags>

Flags:

```
-h          :   This flag will print a help message of how to use the program.
-n <int>    :   The program is set to receive this many packets. This flag is
                optional, if not given the default value is 10 packets.
-l <int>    :   The program is set to receive packets of the given size in bytes.
                This flag is optional, if not given the default is 10 bytes.
-d <int>    :   This flag activates debug printing. The program distinguish between
                1 and >1.
-f <string filename>
            :   If this flag is given the program prints a log to a file given by
                the string. If the file already exists the program terminates after
                given a message.
```


http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

The server listens on port 4445 for UPD packets sent by the client. The server terminates after it has received the amount of packets it should or if it receives a packet with a sequence number exceeding the number of the amount of packets to receive. After the termination of the receive loop the server prints; information about how many packets it received and how many that was lost, information of how long time it tock from the first packet to the last, given in milliseconds and information of the throughput in Kbits/s. If the debug flag is set you could get a printout of every received packet as they arrive. You could also get information about which packets that got through printed, with received packets illustrated with a '*' and non-received illustrated with a '.'.

The log file begins with a full date, which is written when the file is created. Followed by a blank line and then the flags given to the program. If the program in started with the debug flag set, all received packets with sequence number and time offset, separated with a ':', is written to the log. After the loop that receives packets, the program prints a compilation that tells how many packets that where received, for how long time in milliseconds the program received packets and calculated throughput value in Kbits/s.

8.3 The timetable

The work started in the middle of August 2000. The project definition was initially quite diffuse. The first plan was to use an Agilent Viridia monitor and extract data from it. This data would then be sent on an Ethernet network to a computer, which would present it to the screen. The first step was to read about the Viridia monitor. After some contacts with Agilent in London it became clear that it would be impossible to extract data from this monitor. This forced us to look for alternatives. The work continued with searching for companies that delivered similar systems that could be used. After some research I found that Ortivus had a system that did most of the job. In discussion with Ortivus we found out that they are testing their equipment with Bluetooth. This led to an investigation of the wireless aspects of telemedicine.

In the middle of October the Bluetooth cards arrived. The first test was conducted at a research group meeting where an experimental communication via Microsoft Netmeeting over a Bluetooth link was conducted. We quickly drew the conclusion that the transfer-range was greater than expected and that the software installation was easy. November was mostly spent on learning Linux since the plan was to conduct the upcoming tests in a Linux environment. I installed and configured one computer, read some books on the topic, learned to work with the ORiNOCO Wireless LAN card and configure the PCMCIA settings under Linux. It would later come clear to us that it would take long time before there would be any commercial Bluetooth drivers for Linux available. As I waited for Linux drivers I started to look at packet capture and packet generation programs that could be used in the measurements tests. Some of them are described below in the appendix A.

After some discussion about the different aspects and problems with the packet generation programs I started developing a Client-Server software. The reason why I found this necessary was that the favorite packet generation and sniffing program, TTCP, did not provide enough information and that the termination of the TTCP program is done with four sequential stop-packets. These four packets where often lost when the connection became weak which caused the program to lock. Other packet generation programs had other disadvantages that also made them unsuitable. The work on the Client-Server program took the main part of December and beginning of January.

February to April was spent making measurements and tests, which took much more time than expected due to the observation of unexpected behaviors and the amount of time each test-run took. May and June were spent on writing the thesis report.

The initial goal was to be done in time for Christmas 2000.

http://www.it.uu.se/research/reports/2001-028/
Copyright © 2001, Magnus Berggren

Recent technical reports from the Department of Information Technology

2001-027 K. Mahata, T. Söderström, M. Mossberg, L. Hillström and S. Mousavi: *On the use of flexural wave propagation experiments for identification of complex modulus*

2001-026 Jan Gulliksen and Inger Boivie: *Usability Throughout the Entire Software Development Lifecycle - A Summary of the INTERACT 2001 Workshop*

2001-025 Emmanuel Beffara and Sergei Vorobyov: *Is Randomized Gurvich-Karzanov-Khachiyan's Algorithm for Parity Games Polynomial?*

2001-024 Larisa Beilina, Klas Samuelsson and Krister Åhlander: *A hybrid method for the wave equation*

2001-023 Pierre Flener, Alan Frisch, Brahim Hnich, Zeynep Kiziltan, Ian Miguel and Toby Walsh: *Matrix Modelling*

2001-022 Pierre Flener, Alan Frisch, Brahim Hnich, Zeynep Kiziltan, Ian Miguel, Justin Pearson and Toby Walsh: *Symmetry in Matrix Models*

2001-021 Inger Boivie: *Usability and Design Decisions in Software Development*

2001-020 Emmanuel Beffara and Sergei Vorobyov: *Adapting Gurvich-Karzanov-Khachiyan's Algorithm for Parity Games: Implementation and Experimentation*

2001-019 Wendy Kress and Jonas Nilsson: *Boundary conditions and estimates for the linearized Navier-Stokes equations on staggered grids*

2001-018 Emad Abd-Elrady: *An adaptive grid point RPEM algorithm for harmonic signal modeling*

2001-017 Henrik Björklund, Viktor Petersson and Sergei Vorobyov: *Experiments with Iterative Improvement Algorithms on Completely Unimodal Hypercubes*

2001-016 Robert Stjernström: *User-Centred Design of a Train Driver Display*

2001-015 Magnus Svärd: *On coordinate transformations for summation-by-parts operators*

2001-014 Paul Pettersson and Sergio Yovine: *Workshop on Real-Time Tools (Proceedings)*.

2001-013 Eva Olsson, Lena Kecklund, Michael Ingre and Anders Jansson: *Lokförarens informationsmiljö och ATC. Ett användarperspektiv*

2001-012 Friedhelm Stappert, Andreas Ermedahl and Jakob Engblom: *Efficient Longest Executable Path Search for Programs with Complex Flows and Pipeline Effects*

2001-011 Torsten Söderström, Umberto Soverini and Kaushik Mahata: *Perspectives on errors-in-variables estimation for dynamic systems*

2001-010 Thiemo Voigt and Per Gunningberg: *Dealing with Memory-Intensive Web Requests*

2001-009 Johan Bengtsson: *Reducing Memory Usage in Symbolic State-Space Exploration for Timed Systems*

2001-008 Viktor Petersson and Sergei Vorobyov: *Parity Games: Interior-Point Approach*

2001-007 Sven Montan: *Validation of Cycle-Accurate CPU Simulators against Real Hardware*

2001-006 Erik K. Larsson and Torsten Söderström: *Identification of Continuous-Time AR Processes by Using Limiting Properties of Sampled Systems*

2001-005 Yordanos G. Beyene: *Firewalls in Linux: Principles and Implementation*

2001-004 Bernhard Müller and H.C. Yee: *High Order Numerical Simulation of Sound Generated by the Kirchhoff Vortex*

2001-003 Karl Hörnell and Per Lötstedt: *Adaptive Iteration to Steady State of Flow Problems*

2001-002 Johan Steensland: *Dynamic Structured Grid Hierarchy Partitioners Using Inverse Space-Filling Curves*

2001-001 Viktor Petersson and Sergei Vorobyov: *A Randomized Subexponential Algorithm for Parity Games*

UPPSALA UNIVERSITY

November 2001
ISSN 1404-3203
http://www.it.uu.se/