AN IMAGE CAPTURE SYSTEM FOR USE IN TELEHEALTH

Ed Wyn Lim

B.Eng (Hons) in Electrical and Electronic Engineering, Adelaide University

A thesis submitted in fulfilment
of the requirements for the degree of
Master of Engineering

from

The University of New South Wales
School of Electrical Engineering and Telecommunications
Sydney, Australia

31 August 2006

Acknowledgements

I am indebted to my supervisor, Professor Branko Celler, for his guidance, encouragement and financial support during my studies. Also, I am grateful to my co-supervisor, Associate Professor David Taubman, for his technical guidance especially involving JPEG2000 and Kakadu.

I would like to thank Associate Professor Tim Hesketh for his guidance on my project; the staff of the Biomedical Systems Laboratory, especially Hugh Garsden, Jim Basilakis, Khang Huynh and Neil Mackenzie for providing me with much aid during my project.

I would like to thank the following for their friendship and support: Calvin Kwok, Desmond Heng, Siti Muhammad, Rudi Rudi, Teddy Cheng, Leroy Chan, Greg Chan, Einly Lim, Alka Kaushik and Lucy Wang.

I would like to express my special gratitude to my parents and family, for their continuous support in every way possible and patience during my project.

Abstract

The growth of community health and nursing home facilities has led to the demand in functional and effective mechanisms for the management of chronic illness in such facilities. The UNSW Biomedical Systems Laboratory Telemedcare system is one such system currently in deployment. Its modularity lends itself to the addition of different health modalities to tailor to the myriad needs of different patients.

In this project, a practical imaging system for telehealth is developed for and in conjunction with the Telemedcare system to address the need for an imaging modality. The project involves developing an image capture system for use at the patient workstation, a mechanism for delivery and storage of the captured images to a central server, and a mechanism for image retrieval at the client end via a web interface. Several technologies are considered during the course of this project.

To interface with the range of consumer digital cameras from different vendors, the Picture Transfer Protocol (PTP) is used. The existing Telemedcare synchronization and database infrastructure is leveraged to deliver and store the captured images. The JPEG2000 imaging format allows for efficient compression, scalability and codestream flexibility amongst its many benefits. Utilizing Kakadu software by Taubman, the captured images are converted to JPEG2000, while a viewer application and image proxy enable efficient browsing of the images via the JPEG2000 Interactive Protocol.

A full clinical trial could not be carried out. However an evaluation panel of various health specialists finds the system easy to use, and feels that it would be beneficial for use in healthcare in the community and in remote areas.

Contents

List of Figures

List of Tables

Chapter 1

Description of Imaging System

1.1 Telemedcare

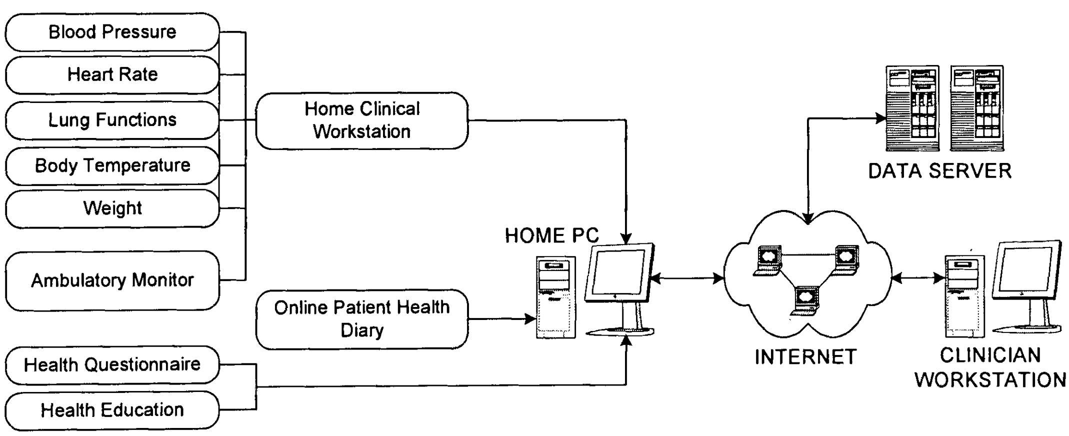

Figure 1.1: Structure of Telemedcare system

The in-house home telecare system, or Telemedcare (previously known as Hetemis [10], [11], [12]), provides a home patient with self-administered functional health status measures in the form of questionnaires, remote monitoring of parameters of daily changes to patient status, and vital signs monitoring. Acquisition of biomedical data or questionnaire responses is done via the Internet. Data are stored in a database and later analyzed for longitudinal trends. These data are available on request to the client (usually a general practitioner) on a web browser.

Figure 1.1 shows the structure of the Telemedcare system in its current version. The patient is provided with a clinical workstation or more formally known as the Home Desktop. It is modular and currently includes the following modules: electrocardiogram (ECG), blood pressure, spirometer, and blood oximetry modules; a low power radio-frequency interface to an ambulatory patient worn device for telephone voice connection on emergency button press and triaxial accelerometer; accurate patient weight and temperature measuring modules; ambient temperature, light and humidity measuring modules. The clinical workstation may also be connected to a web television. For home telecare, the minimization of cost is important, and the home telecare system achieves this by emphasizing modularity in its design, so that the modules required can be tailored to suit individual patients' needs. The patient user interfaces are designed to be simple and intuitive, so that users of the system need not have had any prior experience using computers.

In addition, the home clinical workstation has the following features. It has highly secure remote access to a central server from anywhere via an internet connection. A doctor or clinical team manages the patient with respect to medications and clinical measurement schedules, access to validated health information, and schedules for retrieving patient data from the home. The doctor or clinical team has access both to patient longitudinal records and the original clinical data recorded. The home telecare system can be controlled and managed remotely, facilitating its use in rural and remote areas.

Based on case studies and clinical trials, the Telemedcare system has been shown to be functional and effective. It has been deployed in actual home settings and has been found to be a useful tool for the management of chronic illnesses like congestive heart failure (CHF) and chronic obstructive pulmonary disease (COPD) [12]. One such trial was conducted in 2002 by the Commonwealth Department of Health and Aging. In this trial, 22 patients aged 58 to 82 from Sydney and Wagga Wagga in rural New South Wales, were monitored using the home clinical workstation for a period of 6 to 12 months. These patients had primary diagnoses of CHF and COPD and were hospitalized for the conditions within the last six months. The patients' primary care physicians were recruited for this trial, and both patients and physicians were administered with initial, follow-up and final evaluation questionnaires to be submitted for the trial. 89% of participants were satisfied with the system while no one was dissatisfied (11% were undecided). The same percentage of physicians felt that the system could play an important role in the management of patients' health. The design of Telemedcare is driven by specific clinical needs, and the successes of the trials are a validation of its relevance to those needs. By satisfying the clinical needs of patients with chronic illness without leaving their homes and reducing the frequency of visits to a medical institution, Telemedcare has the potential to be a cost-effective and convenient alternative of health care provision.

The Australian Institute of Health and Welfare (AIHW) released a report, Residential Aged Care in Australia 2003-04: A Statistical Overview [14], about the current developments and growth of residential care facilities in Australia. According to the report, the number of residential aged care places in Australia and the ratio of places to people aged 70 years and over increased in the years 2003-04. The report also showed that there were 156,580 residential aged care places at the end of the 2003-04 financial year, which was an increase from 151,181 places the previous year, even though the number of mainstream aged care services actually declined slightly from 3015 at the end of the 1997-98 financial year. There were 84.2 residential places for each 1,000 people aged 70 years and over, compared to 82.8 places per 1,000 a year earlier. The proportion of residents classified as 'high care' rose from 61% in 1999 to 66% in 2004. The average age of residents also increased during this period from 49% to 51%.

The statistics from the AIHW report demonstrate that the demand for residential care facilities is increasing. At the same time this increasing demand imposes a greater burden on the nursing community involved in the care of residents in these facilities. Consequently the need for more cost-effective solutions for health care provision and health monitoring becomes more apparent.

1.2 Imaging Modalities for Telehealth Systems

1.2.1 Common modalities

Telemedcare does not currently have any facility for imaging incorporated in it. Having a system for imaging will greatly extend the functionality of Telemedcare as it will allow the general practitioner to have access to a range of visual health status measures in addition to the existing measures. The patient will then be able to report external afflictions such as wounds, skin lesions and rashes, as well as provide the general practitioner with an accurate representation of any part of the patient's body if desired.

The imaging modalities that are most commonly used in telemedicine are the 'store-and-forward system' modality [25], [26], [29], where images are saved locally, then forwarded by email for instance to the recipient, and the real-time video/web conferencing modality [13], [25], [29].

In [29], the advantages of both modalities are listed. The main advantages of the store-and-forward modality are that it is a cheap and effective means of providing diagnosis to the patient, it has adequate diagnostic accuracy, and the physician may review large numbers of images at his/her own convenience. Real-time conferencing obviously allows interactive consultation between the physician and patient. Real-time conferencing also provides a more accurate mechanism for diagnosis and patient management compared with store-and-forward techniques, and can make available a greater amount of clinical information. However, both modalities also have their disadvantages. The physician is unable to communicate directly with the patient with the store-and-forward modality. On the other hand, it may be difficult to synchronize the patient and physician for a conferencing session. Both parties may be frustrated by the technical issues involved in setting up the conferencing equipment. Also, a conferencing session may be as time-consuming as a regular physician consultation.

A key finding of [29] is that, for teledermatology, the store-and-forward modality is much more economical and cost-effective compared to real-time conferencing, but by permitting interaction and consultation between the patient and physician, real-time conferencing is more clinically effective.

The communications infrastructure is still dominated by conventional telephone lines with a typical bandwidth of 33.6-56kbps, but digital lines with bandwidths ranging from 128-512kbps are becoming more commonplace. However, the availability of digital lines is usually confined to urban or more densely populated areas where it is more cost-effective to lay out digital lines. Areas which are not serviced by a digital line are thus restricted by the slow (and perhaps) poor quality telephone line. In this way, for both modalities, their use over a network with limited bandwidth hampers their effectiveness. In the former, the GP may need to wait for a long time for large images to arrive. For example, it would take 1.7 hours for a store-and-forward system to transmit an image of size 42MB over a telephone line with a nominal bandwidth of 56kbps. In the latter, the conferencing session may experience high delay and latency. On a telephone line, motion handling is poor and motion artefacts are increased [29]. Digital lines are much preferred for real-time conferencing systems. In addition, conferencing systems tend to incur higher implementation costs. These could include the cost of the digital link and the conferencing equipment themselves. The additional cost is particularly acute if conferencing is to be employed over short distances. With expensive conferencing equipment, equipment failure becomes a major concern when determining the cost-effectiveness of this modality.

The network links between the home workstation and central server, and central server and client workstation on the client side, will likely be of low bandwidth, e.g. a conventional telephone line with a bandwidth of 56kbps. From the previous discussion, both the store-and-forward and real-time conferencing modalities are likely to suffer from delay and high latency if put into operation. This may then lead to unfavourable outcomes such as low user satisfaction.

1.2.2 JPEG2000 and Progressive Transmission

This project attempts to utilize progressive image transmission as an imaging modality for Telemedcare. We use the JPEG2000 imaging standard which allows for efficient compression, scalability and codestream flexibility amongst its many benefits [2], [7]. The standard allows progressive transmission of compressed images based on precinct, component, resolution or quality. The standard also defines the JPEG2000 interactive protocol (JPIP), which is based on HTTP, for efficient serving of images over data networks [1], [5], [9].

With JPIP we can efficiently browse (or serve) large images remotely, and with a judicious choice of progression we can browse images with good quality (in the distortion sense) even over low datarate networks. If an image is coded with several quality layers, receiving the initial layer causes the client to be presented with a low quality image at a certain resolution initially, and then receiving the following layers improves the quality of the image gradually. A good and perhaps useful representation of the original image is possible even with a small number of layers.

Figures 1.2 to 1.5 show the image of an eyeball in its entirety, which has been magnified. The image was progressively transmitted at an artificially throttled bandwidth of 0.5kB/s, and the figures show the image at various stages of transmission.

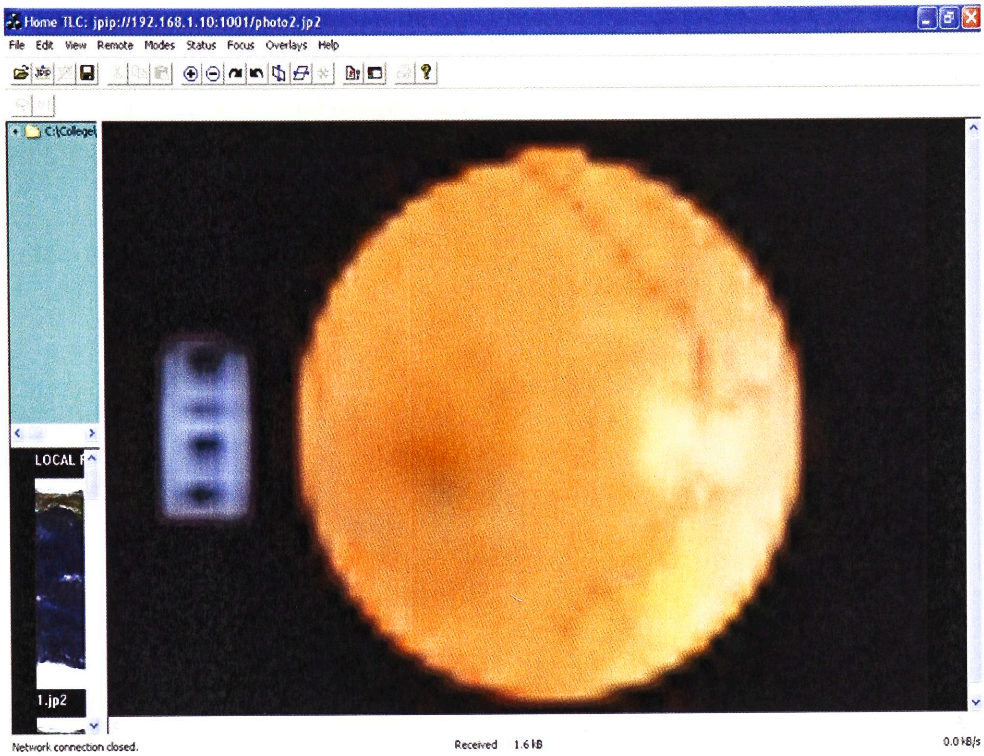

Figure 1.2: Image of eyeball at 1.6kB received bytes

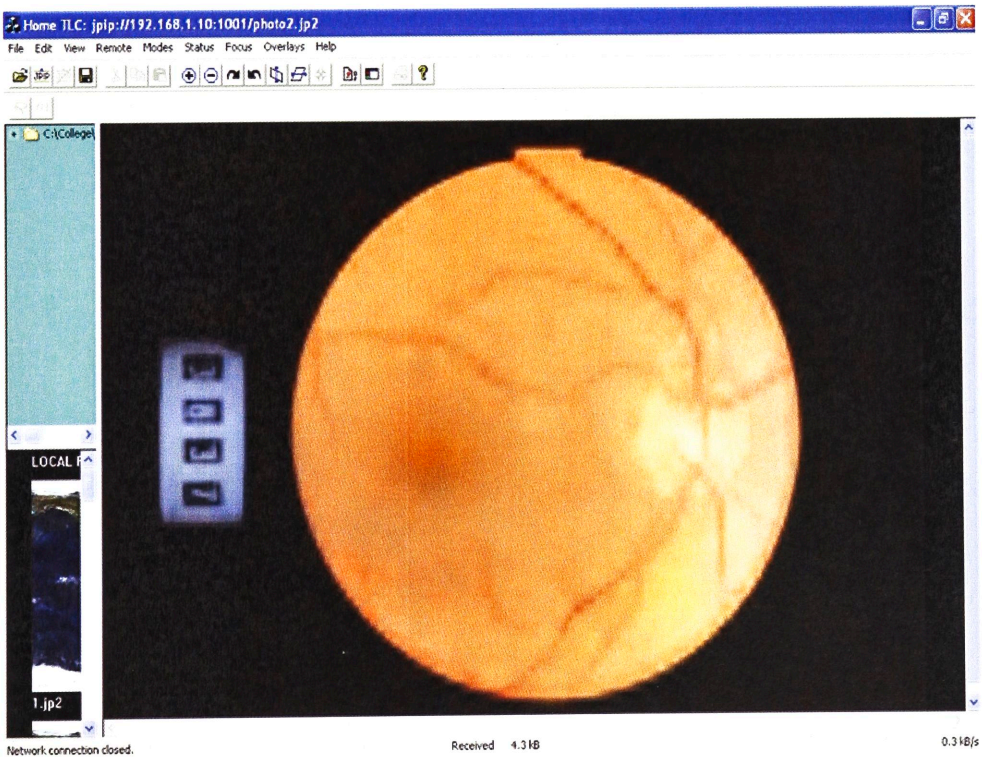

Figure 1.3: Image of eyeball at 4.3kB received bytes

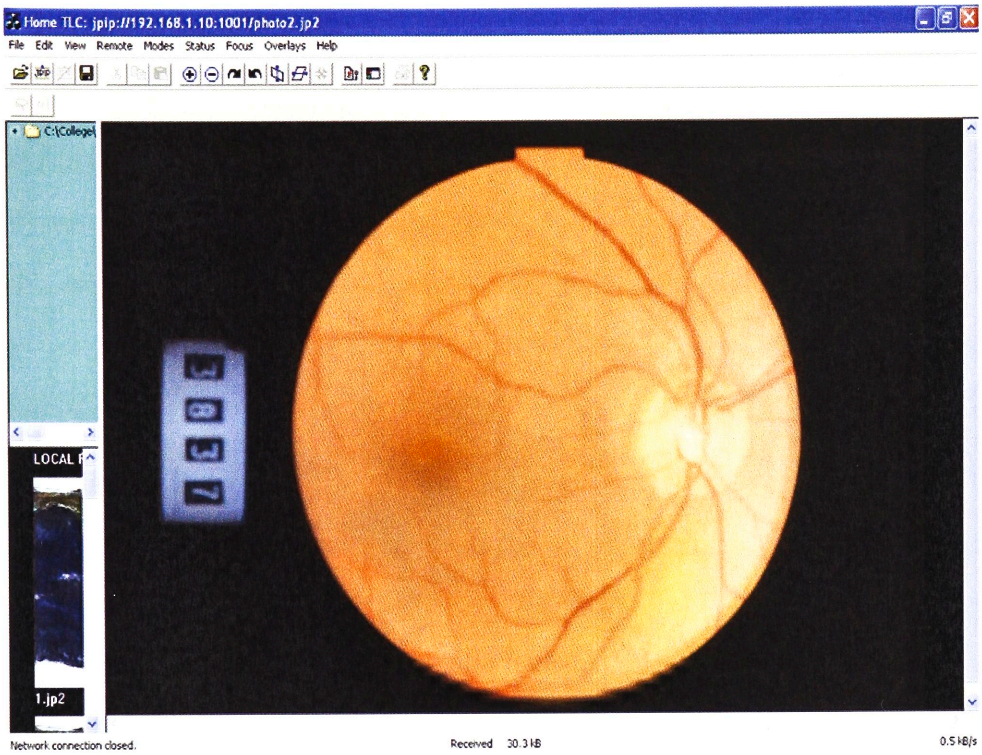

Figure 1.4: Image of eyeball at 30.3kB received bytes

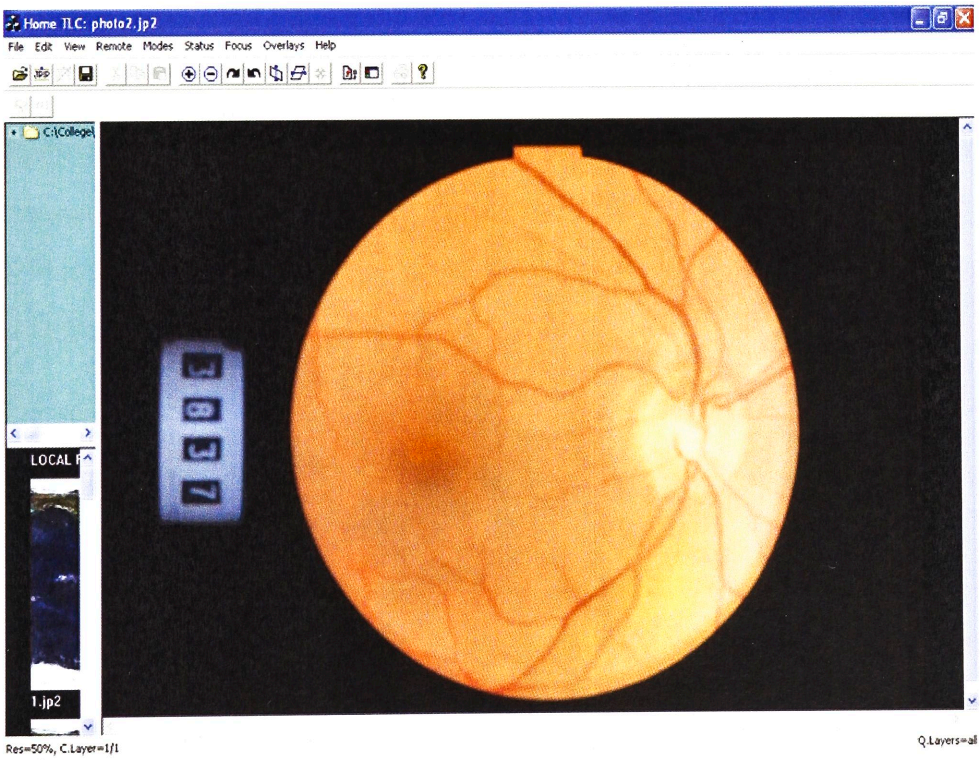

Figure 1.5: Image of eyeball at full size (2.62MB)

When 1.6kBs are received (Figure 1.2), the eyeball is extremely blurry but its outline and the largest blood vessels are roughly visible. When 4.3kBs are received (Figure 1.3), the image is still blur, but the iris is now visible and more blood vessels show up; in addition the numbers on the measuring device on the left can now be read. When 30.3kBs are received (Figure 1.4), the blood vessels have been greatly enhanced, and the eyeball is much sharper overall save for the bottom area. At the full image size of 2.62MBs (Figure 1.5), the eyeball image is virtually indistinguishable from Figure 1.4, except for the bottom area which has now been brought to full quality.

In contrast, other common image file formats that may possibly be used in web or interactive environments such as JPEG, GIF, TIFF and PNG for instance, typically appear on screen in raster fashion. In the case of JPEG, an extension was added to the standard to allow JPEG images to be coded in such a way as to allow such images to be browsed progressively in the same manner as a JPEG2000 image (hence the term progressive-JPEG for JPEG images coded this way). The difference between progressive-JPEG and JPEG2000 is simply that the latter image standard was designed from the ground up with scalability and progression in mind.

There are a number of commercial products which utilize JPEG2000 and progressive transmission in healthcare, for instance the Aware JPEG2000 software for medical imaging [15]. In the literature, telehealth projects using JPEG2000 are somewhat harder to find.

1.2.3 Existing telehealth systems

One project which does use JPEG2000 for teleradiology is the IDEAS project [16]. The IDEAS in e-Health project presents a general architecture that aims to provide a common platform which supports a large variety of telehealth services such as telemedicine, e-care, tele-assistance and teleradiology, oriented to general medical assistance. IDEAS is based on a web-based Application Service Provider (ASP) platform. The ASP server is implemented using a distributed architecture that enables multiple servers to share the workload and scale to larger demands. The ASP uses standard components and provides access to the applications with just a web-browser and Java virtual machine, and a reliable Internet connection to access stored resources. The project demonstrates that the ASP model can work for several telemedicine applications enabling hospital users to downsize the cost of installing, managing and upgrading computing resources. Data can be also securely available from different sites, wherever they are needed. For teleradiology, medical images in DICOM 3.0 or ACR/NEMA formats are encoded using segmented JPEG2000 lossless compression, thus preserving their diagnostic value. Image transmission is performed as a background process using progressive transmission and visualization. The image diagnosis application is written in Java and can be executed on any computer. Diagnostic reports can be attached and stored securely on the ASP for further consultation from the hospital. A gateway to high-performance computing is provided for 3D volume operations, such as segmentation and projection. External storage of the data and provision of computer-based diagnosis tools enable radiologists to provide complete diagnosis from any secure Internet connection, just requiring a web browser. Progressive downloading of images at lossless resolution in the background, allows the user to avoid waiting until the last image is available. For home telecare, a home telecare application has been developed with Microsoft '.NET' technologies and connections are made secure through Secure Sockets Layer and X.509 certificates. Connectivity is achieved through the use of XML, IP, GPRS, IEEE802.11b, combining emerging wireless technologies with the existing fixed network structures.

The system in [19] is a calibrated store-and-forward system for use in teledermatology. It uses high-resolution cameras, proper illumination and polarizers to avoid skin reflections, low resolution due to equipment limitations, and noise due to sensor limitations or particles on skin such as hairs. For data transmission, the system considered using the public telephony network, but settled on Integrated Services Digital Networks (ISDN) networks and TCP/IP to accommodate the large amount of data that have to be transmitted reliably and cheaply. As delay can be tolerated, the system uses the store-and-forward paradigm for asynchronous and lower bandwidth transmission. The average time required for the examination of the images and the diagnosis is about six minutes, which covers the examination of images, the patient's file, and other transmitted data. Data throughput is improved by using compression on the images, which take up the most bandwidth among the transmitted data. The type of compression used is a from of wavelet compression, using the CDF-9/7 biorthogonal wavelet.

A web-based medical image archive system in a three-tier, client-server architecture for the storage and retrieval of medical image data, as well as patient information and clinical data, is presented in [17]. The Web-based medical image archive system was designed to meet the need for a central image repository to address questions of stroke pathophysiology and imaging biomarkers in stroke clinical trials by analyzing images obtained from a large number of clinical trials conducted by government, academic and pharmaceutical industry researchers. In the database management-tier, image storage hierarchy is designed to accommodate large binary image data files that the database software can access in parallel. In the middle-tier, a commercial Enterprise Java Bean server and secure Web server manages user access to the image database system. User-friendly web-interfaces and applet tools are provided in the client-tier for easy access to the image archive system over the Internet. Benchmark test results show that the three-tier image archive system yields fast system response time for uploading, downloading, and querying the image database.

The system in [18] tests a simple, fast, and inexpensive method for practical transmission of images for diagnosis using a digital camera and the Internet. Using a commercial digital camera mounted with a phototube adapter to a light microscope (6 images per case on average), 2210 digital images (310 Mb) from 347 cases of gastrointestinal, lung, and uterus specimens were captured. Each image, stored in medium compression JPEG (Joint Photographers Experts Group) format with 1024 3 768 pixel resolution, required approximately 5 seconds to capture after the case had been reviewed and appropriate fields for imaging selected (30 seconds per case on average). The images were transmitted from Samsung Medical Center, Seoul, to Korea University Hospital, Seoul, and John Hunter Hospital, Newcastle, Australia. Transmission was 100% successful with a total upload time of 3 hours for 310 MB of data (31 seconds per case on average). The images were downloaded in 2 hours and viewed on a 17-inch color monitor with a maximal resolution of 1280 3 1024 pixels. Telepathology diagnoses were made with 95% and 97% concurrence by two pathologists at Korea University Hospital and John Hunter Hospital, respectively.

1.2.4 Use of COTS imaging devices in Telemedcare

The clinical trial of Krupinski et. al.

The use of commercial-off-the-shelf (COTS) digital cameras in telemedicine is not new. In [27], such cameras used in the context of teledermatology were found to provide clinicians with high-quality images suitable for diagnosis, with outcomes comparing well with in-person consultations.

In the clinical trial, one of three certified dermatologists performed a general examination on the patient. There were 308 patients who participated in this trial. The patient's skin ailments were then captured using a digital camera. In this trial a Canon Powershot 600 camera with a spatial resolution of 832 x 608 and 24-bit resolution was used. For each skin ailment, images were taken first of the general region, followed by a close-up of the affected area, with distances ranging from 5 to 20 cm. For each ailment, 5 images were taken in total.

The images were divided into separate cases and the patient's case history from his/her medical record was attached to the image case file. A brief patient history derived from the patient's medical record was inserted into each case file. Care was taken to include only patient history information and exclude any reference to actual diagnosis, patient name, or the identity of which dermatologist had examined the patient. A total of 308 cases for each patient were obtained, these were randomized and numbered for presentation order during the study. The same three dermatologists reviewed the digital images after a lapse of two months, to encourage forgetting the previous diagnoses that had been made. Patient history was noted prior to viewing the images for each case.

A single most likely diagnosis was rendered with an associated degree of confidence ("very definite," "definite," probable," "possible") for each case. Subjective ratings of image sharpness and color were given using a 4-level scale (excellent, good, fair, poor). Viewing time was recorded using a stopwatch, beginning from when the first image of a case appeared on the computer monitor, and ending when the dermatologist rendered a diagnosis. The dermatologists were restricted to 1-hour viewing sessions to minimize fatigue.

The results of the clinical trial are summarized as follows. There was 83% concordance between in-person versus digital photo diagnoses. Intradermatologist concordance averaged 84%, and interdermatologist concordance averaged 81%. Decision confidence was rated as "very definite" to "definite" 62% of the time. Concordance with biopsy results was achieved in 76% of the cases. Image sharpness and color quality were rated "good" to "excellent" 83% and 93% of the time, respectively. It is reasonable to say that, since this paper and its corresponding clinical trial were produced in 1999, and in the 6 intervening years the technology of digital cameras has improved tremendously, if a similar trial were to be conducted contemporarily, the trial would produce similar if not superior results as that obtained from [27]. Indeed, such a trial was conducted in 2006 and described by Murphy et. al. in [28].

Recommended features and usage methods of digital cameras

In [29], some recommended features when considering a digital camera for use in a teledermatology setting are mentioned as follows:

- a minimum resolution of 1024×768;
- white load balance;
- USB connectivity, which allows direct image download from camera to computer. This can be extended to wireless and Firewire connectivity;
- memory card storage (4MB - 128MB); memory cards of up to 2GB are commonly available nowadays;
- brightness or EV compensation (ideal -2.0 to +2.0);
- aspherical glass lens;
- optical zoom;
- red eye compensation;

In [31], some recommended usage methods for a digital camera for use in a general telemedicine system are mentioned as follows:

- composition and lighting are the keys to better photographs
- the background should be chosen carefully; for instance, a plain wall should be used and the window should not be included
- the image should be limited to the particular area of concern
- clothing and jewellery should not obscure the image

- it should be possible to identify the area of interest, and there should be sufficient detail for the physician to make a diagnosis
- the angle of view should be carefully considered (e.g. consider if the area of interest is raised or flat)
- a general view should be provided to show the location and scale of the condition, and a close-up view should be provided to show detail
- the subject should check that it is in focus.

Chapter 2

JPIP and JPEG2000

2.1 JPEG2000

A JPEG2000 image can be divided into tiles which are compressed independently. Each tile in the image is decomposed into spatial frequency subbands using the Discrete Wavelet Transform (DWT). Each subband is then partitioned into small blocks known as code-blocks. Each code-block is then coded independently according to the Embedded Block Coding With Optimized Truncation (EBCOT) paradigm. Each resolution in the DWT is partitioned into precincts, with every code-block belonging to exactly one precinct.

In order to allow for progressive improvement of image quality, JPEG2000 introduces the concept of quality layers. Each layer is formed by sample contributions from all code-blocks in the image, although some code-block contributions may actually be empty. In every precinct, the quality layer contributions from the code-blocks belonging to the precinct form a packet. A packet consists of a header containing information regarding the contributions from each code-block, followed by a body containing the code-block contributions themselves. Although JPEG2000 also defines three other

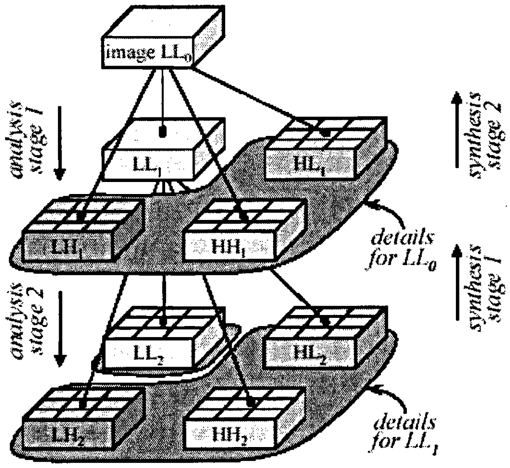

Figure 2.1: Image subbands and codeblocks (*image courtesy of [1]*)

types of progressions (progression by resolution, progression by component, and progression by spatial content), we will be only concerned with quality layer progression here. Referring back to the eyeball images in section 1.2.2, the most natural form of viewing employed by a clinician would be to view an image at a fixed resolution (or image dimensions), and allow the quality of the image to be gradually increased (or alternatively allow the distortion of the image to be gradually reduced) at this fixed resolution.

2.2 JPEG2000 and metadata

In the context of JPEG2000, the term metadata has two separate meanings. Firstly, metadata can assume its customary meaning of data about data. Secondly, it means a collection of boxes of a JPEG2000 file format.

2.2.1 Metadata as a collection of JPEG2000 boxes

The core JP2 file format, ".jp2", and the extension format, ".jpx", are comprised of *boxes*, each of which contains different portions of the image codestream. In general, boxes are free to be placed anywhere in the file, with some minor restrictions. Boxes consist of the 4-byte box length, followed by the 4-byte box type, and then the box contents. The box length indicates whether the box is in fact a regular box or a super-box, where in the latter case an extended length will be read before the box contents. Figure 2.2 shows some JPEG2000 boxes of interest.

The JPEG2000 signature box identifies the file as being a member of the JPEG2000 file family; one and only one must appear first in the list of boxes. The next box is the file type box, which identifies the family type that the file is. For instance "jp2" signifies the core format, and "jpx" signifies the extended format. Following the file type box is the JPEG2000 header box. This box is a little more interesting in that it is a superbox, containing no unique data of its own except for other sub-boxes. Some sub-boxes that are contained in the header box include the image header box, the bits per component box and the palette box. This initial set of boxes customarily comes first in any given JPEG2000 file and normally contains the information required for proper decoding and rendering of imagery data (the extended "jpx" format defines several other required boxes). What follows is the contiguous codestream box, which contains the actual JPEG2000 codestream, and more than one codestream box may appear in the extended file format.

The boxes which come after the contiguous codestream box are considered auxiliary because they are not required for rendering the imagery. These boxes represent metadata in the sense of data description, as is explained in the next subsection. Such boxes include the intellectual property rights boxes, the XML box, the UUID boxes and the UUID information boxes. Metadata boxes can actually appear anywhere in the file, but must appear after the JPEG2000 header box.

This interpretation of metadata as a collection of boxes will be most heavily relied upon in the section on JPIP (Section 2.3).

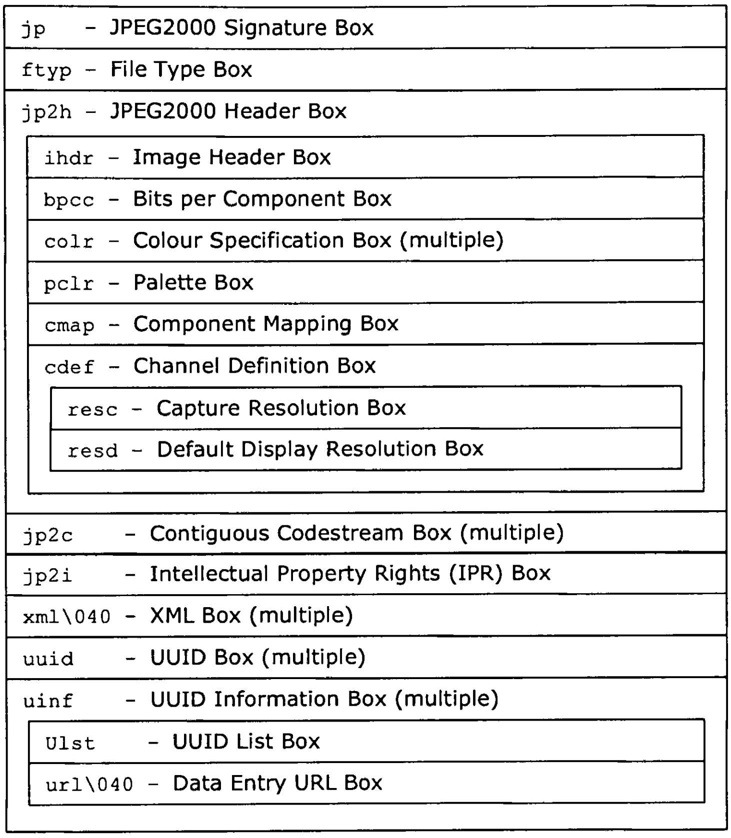

Figure 2.2: Sample JP2 Box Structure

2.2.2 Metadata as auxiliary data about images

The second and more common usage of the term metadata is data which describe other data. Thus image metadata are data that are used to describe the contents of the image. For instance, an image of a scene at the beach may contain metadata such as: the name of the photographer; a description of the scene in the image; a description of the color composition of the image; the camera settings used to capture the image; the location of the scene. If the image is in digital form, the processing history of the image could be an additional form of metadata associated with it.

As described in the previous subsection, the JPEG2000 boxes which are designed to encapsulate metadata are the IPR, XML and UUID boxes. In this project our focus will mostly be on the XML boxes, as JPEG2000 utilizes XML as the preferred document format to store image metadata. XML is also utilized by the Telemedcare system in a different way, which is mentioned briefly in section 3.2.1.

XML and metadata

In this section we provide a concise discussion on XML and its benefits to document storage and transmission. XML stands for eXtensible Markup Language, and is a standard for document markup. From [33], XML defines a generic syntax to mark up data with human-readable tags, and it provides a standard for computer documents which are flexible enough to be customized for various areas and applications.

In an XML document, strings of text are surrounded by tags or markup which are known as elements. XML is flexible because the set of tags which are to be used by an application depend on the application developer, and XML itself does not specify any set of tags to be used. XML however does specify a consistent syntax for the forms of tags and markup that any XML developer has to strictly adhere to. For example, tags must have names which must not begin with numbers and must not contain characters such as whitespace, and '<' and '>' which are used to delimit the tags themselves. In another example, names must not begin with the sequence 'xml' as it is reserved by XML. XML documents which are produced to follow such fairly rigid grammar and syntax are called well-formed XML. [33] makes the distinction that XML is a 'structural and semantic' language, and not a 'presentation' language like HTML with which XML bears some resemblance.

The extensibility of XML necessitates a mechanism to differentiate the tags used in one set of XML tags from another, since multiple sets may be contained in a document. For instance, the tag '<table>' can be used to describe either a piece of furniture or a structure used for web display. This mechanism is defined in XML as namespaces, where different namespaces are identified by a unique Universal Resource Identifier (URI) such as a HTTP hyperlink. So, the two instances of the element table above can be differentiated by adding a prefix linked to a unique URI, say <furniture:table> and <display:table>. The set of XML tags themselves that are allowed to be used in an application are defined in a schema. The XML specification defines the document type definition (DTD) and the World Wide Web Consortium (W3C) have produced a schema of their own, the XML Schema Language. Thus, when an XML document is compared with its corresponding schema and the tags and syntax within the document are found to adhere to the schema, the document is then considered valid. XML however does not enforce validity, only well-formedness; the enforcement of validity is an application issue.

Other constructs used to interact with XML that we are concerned with are: XPath, a syntax for searching tags in an XML document; and eXtensible Stylesheet Transforms (XSLT), a method of transforming the structure of tags in a document especially to aid presentation.

XML Metadata and JPEG2000

Part 2 of the JPEG2000 standard [8] provides extensions to the core portion of Part 1. This includes a comprehensive set of optional metadata elements that may be written into XML boxes in an extended JPX file. The image metadata that can be embedded in a file are conceptually divided into the following groups:

- Fundamental data types and elements: Primitive data types that are to be used in the definitions of elements from the other groups are defined here for ease of reference. These include the customary types like integer, degree, rational and double as well as complex types like person, organization, phone and email types for identification, and GPS, point, region and position for location.

- Image creation metadata: This group contains elements that permit the description of how the image came to be. As expected camera and scanner parameters and settings are well-represented by elements that allow technical aspects of the image acquisition process to be recorded. This group also includes two elements for print and film media.
- Content description: Elements from this group allow the description of events, locations, scenes and persons depicted in the image.
- Metadata history: All previous editing and version changes to the images, whether to the image data or its metadata, may be recorded using elements from this group.

Although metadata concerning IPR are contained in a separate box, the section on metadata in Part 2 of the standard also defines the IPR elements as they are essentially XML elements as well.

XML metadata which are optionally embedded along with a JPEG2000 image can be tested for validity against either a DTD or a W3C Schema as defined by Part 2 of the standard. The URL which references the DTD is "http://www.jpeg.org/metadata/-15444-2.dtd", while the URL which references the Schema is "http://www.jpeg.org/-metadata/15444-2.xsd". The standard also defines a URI for a namespace that can be used to prefix JPEG2000 XML elements, and it is given by "http://www.jpeg.org/jpx". In this project the focus will be on the use of Schemas in preferences to DTDs.

The following is a sample of image metadata that may appear in a JPEG2000 file [32]:

```
<?xml version="1.0" encoding="UTF-8"?>
<CONTENT_DESCRIPTION
```

```
  xmlns="http://www.jpeg.org/jpx/1.0/xml"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.jpeg.org/jpx/1.0/xml
  http://www.jpeg.org/metadata/15444-2.xsd">
  <CAPTION>Houchin Castle</CAPTION>
  <LOCATION>
    <ADDRESS TYPE="Scene Address Location">
      <ADDR_COMP TYPE="Street">
        269 Castle Street
      </ADDR_COMP>
      <ADDR_COMP TYPE="City">Greece</ADDR_COMP>
      <ADDR_COMP TYPE="State">New York</ADDR_COMP>
      <ZIPCODE>14059</ZIPCODE>
      <COUNTRY>US</COUNTRY>
    </ADDRESS>
  </LOCATION>
  <EVENT>
    <EVENT_TYPE>Moving Day</EVENT_TYPE>
    <DESCRIPTION>Scott's new Castle</DESCRIPTION>
    <COMMENT>Color scientists rule</COMMENT>
  </EVENT>
</CONTENT_DESCRIPTION>
```

The first line declares the file to be in XML format, along with the version number and character encoding. Although the latest version of XML is 1.1, version 1.0 still predominates. The root element of the file is 'CONTENT_DESCRIPTION'. Within the root element, the following are declared: a default namespace, the 'xsi' namespace, and the location of the reference Schema. The root element has 3 children: 'CAPTION', 'LOCATION' and 'EVENT'. In turn, the 'LOCATION' and 'EVENT' elements have their respective children, with some appearing on multiple occasions, such as the 'ADDR_COMP' element.

2.3 JPIP

2.3.1 Message header

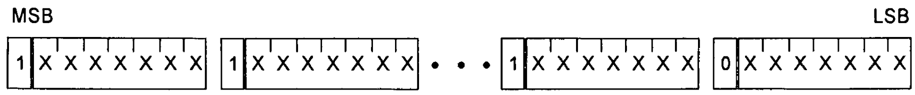

Figure 2.3: JPIP message header (*image courtesy of [9]*)

Figure 2.3 shows the structure of a JPIP message header. Each JPIP message contains a header and a body, and the header identifies the data-bin and byte range represented by the message body. Only one data-bin can be represented by any message. The message header consists of a series of variable-length byte-aligned segments (VBAS), which in turn consist of a series of bytes. These bytes have a most significant bit of 1 except the last byte in the VBAS, signifying that it is indeed the last byte. The message header contains the following VBASs:

$$Bin - ID, Class, CSn, Msg - Offset, Msg - Length, Aux$$

They are respectively, the bin-ID, bin class, code-stream number, data-bin offset, data-bin length, and auxiliary VBASs. The *Class*, *CSn*, and *Aux* VBASs are actually optional and their existence can be determined from other VBASs. Concatenating the least significant 7 bits of each byte in a VBAS produces a bit-stream which is interpreted differently for different VBASs. Message headers can be self-describing, or they can take a dependent form where previous message headers are required to interpret the present message header. Messages with self-describing headers can thus be sent out of order.

Data-bin type	Information	Class
Precinct	Precinct data	0, 1
Tile header	All tile-part headers concatenated in a tile	2
Tile	All tile-parts concatenated in a tile	4, 5
Main header	Main header	6
Metadata	Collection of boxes of a JPEG2000 family file	8

Table 2.1: Data-bin types

2.3.2 Data-bins

JPIP defines the concept of data-bins which make a JPEG2000 image file more a-menable to streaming over a network transport. The image data can be conceptually placed in data-bins. At present the data-bin classes are defined as in Table 2.1.

A JPEG2000 file or codestream can be represented by two media types: *jpp-streams* and *jpt-streams*. Jpt-streams utilize more than one tile to partition the image samples, while jpp-streams are untiled (i.e. the entire image is a single tile). So, data-bins related to tiles are only relevant to jpt-streams, and similarly data-bins related to precincts are only relevant to jpp-streams.

Precinct data-bins provide more flexibility for interactive transmission at the expense of reduced flexibility in determining compression methods and parameters for the image. Henceforth we assume that an image is untiled, and that only precinct data-bins and jpp-streams will be used.

2.3.3 Metadata-bins and placeholders

Besides data-bins, JPIP also defines the use of placeholders to replace the box structures used to encapsulate metadata and codestream data in a JPEG2000 file (Figure 2.4). As noted in Section 2.2 and Table 2.1, metadata can be considered as a collection of boxes of a JPEG2000 family file. Metadata-bin number 0 is conceptually the root metadata-bin, and it may be possible to put all metadata into bin 0. When a

JPEG2000 box is replaced by a placeholder, the placeholder will retain the box header as well as contain the offset to the new metadata-bin which will hold the contents of the box which was replaced. A JPIP server is free to make decisions on how to allocate placeholders in order to best serve a client's request. The server is then able to omit metadata-bins which are not relevant to a client's request. A particularly important point to note is that the metadata-bin structure assigned to a jpp-stream with a unique identifier should not be changed.

A JPEG2000 box may be a superbox, i.e. it contains other JPEG2000 boxes. These boxes may in turn be replaced by a placeholder. However, there is a restriction to this according to [9]: wherever a placeholder replaces a sub-box, a placeholder shall also replace its containing box. This restriction ensures that it is always possible for a client or rendering agent to recover the lengths and locations of the original boxes within the file, even if some of the boxes are not understood by the client.

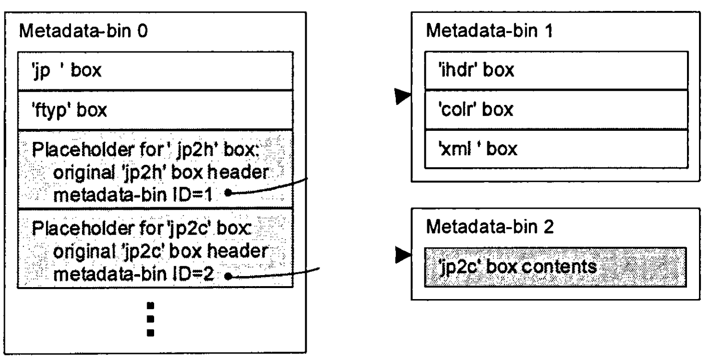

Figure 2.4: JPIP metadata structure and placeholders (*image courtesy of [9]*)

2.3.4 Sessions, channels and cache models

JPIP also defines sessions and caches. Clients may either issue stateless or stateful image requests. If a request is stateful, the server maintains state information about the client capabilities and preferences, and a record of the responses it has already sent to the client within a session, under the assumption that these responses are faithfully cached by the client. If a request is stateless, the request is fully self-contained, and the server does not maintain any state about the client at all. A stateless request tends to be larger as it normally contains information about the client's cache contents, however the server need not allocate extra storage that would otherwise be required for the client information in a stateful session.

In a stateful session, the client responses are recorded by the server by maintaining a cache model. Actually, the server does not strictly have to maintain an accurate cache model, or even a model to begin with. The client has the option to issue statements which directly manipulate the server's cache contents. There are two types of such statements: additive and subtractive. These statements direct the server to alter the composition of the cache model in terms of adding or removing data-bins (or parts of data-bins). Additive statements may be ignored by the server (since it is has no obligation to maintain a model), but subtractive statements must be complied with.

From [9], each session is associated with the following elements:

- One or more logical targets (usually image files), whose content does not change over the session.
- A single image data return type for each logical target associated with the session.
- For each logical target associated with the session, a model of the client's cache contents shall be maintained wherever the data return type is one of "jpp-stream" or "jpt-stream". Note, however, that this model need not perfectly reflect the actual state of the client's cache.
- One or more JPIP channels. Clients may generally open multiple channels within the same session. Each JPIP channel may be associated with a separate underlying transport channel (e.g., a separate TCP connection), although

this might not be the case. Multiple channels allow clients to issue simultaneous requests for multiple image regions, with the expectation that the server will respond to these requests concurrently. Channels also allow for intelligent bandwidth allocation amongst different types of requests either within a single target image or across multiple targets.

- Where multiple channels are associated with the same logical target, the session cache model applies across all channels. Multiple clients may open JPIP channels within the same session, although this might have undesirable side effects if the channels refer to the same logical target.

Also from [9], each channel is associated with the following elements:

- A single logical target (usually an image file).
- A server-assigned identifier that shall be included with each request. JPIP does not define a separate session identifier, since the channel identifier is sufficient to associate the request with its session.
- A record of the client's capabilities and preferences, which may be adjusted through appropriate request fields.
- To the extent that the server queues requests, it should provide a separate queue for each JPIP channel.

2.3.5 Request/response syntax

JPIP request-response messages are basically user-defined extensions of HTTP. A JPIP request is typically framed as a rectangular window into the image, at a particular resolution.

A sample client request looks like this:

```
<< GET /phoenix.jp2?stream=0&len=2000&tid=0&type=jpp-stream&cnew=http
   HTTP/1.1
<< Host: dst-m
```

Each client request begins with a 'GET' or 'POST' statement, and ends with a 'HTTP /version', as with conventional HTTP requests. The request consists of elements which should appear only once within a request.

Some common elements are:

- stream: codestream number
- len: byte limit
- tid: indicates that a new target-ID needs to be assigned
- type: JPEG2000 media type (either jpp-stream or jpt-stream)
- cnew: indicates that a new channel is required

A sample server response looks like this:

```
>> HTTP/1.1 200 OK
>> JPIP-tid: 281B6E135135BBC0BC588452AC9B73C5
>> JPIP-cnew: cid=JPH_033C38BE48115AC9,path=jpip.cgi,transport=http
>> Cache-Control: no-cache
>> Transfer-Encoding: chunked
>> Content-Type: image/jpp-stream
>>
>> 102
-- 258 bytes of binary data
>> 0
>>
```

The first part of the response message is the message header. The 1st line of the header consists of the HTTP version identifier, the status code, and a brief status message. The 1st line is then followed by a series of response fields. Some common fields are:

- JPIP-tid: the target-ID assigned to a JPEG2000 image being requested from the server
- JPIP-cnew: details of the new channel (ID, transport)
- Cache-control: usually set to 'no-cache' to inform normal HTTP proxies not to cache the contents of the JPIP message.
- Transfer-Encoding: Chunked transfer encoding ensures that dynamically produced content can be verified by the recipient by including some state information (chunk length, entity headers).
- Content-type: self-explanatory

The response fields are then followed by the message body. If the encoding is chunked, the message body will be a hex-encoded number indicating the number of bytes being sent, followed by the actual message contents. The fields are terminated by a 'carriage return-line feed' pair. The entire message is terminated by a null (0) terminator.

2.4 Caching With JPIP

Ortiz et. al. [4] proposed modifications to the JPIP specification to exploit existing HTTP proxy infrastructure for interactive browsing of JPEG2000 images. An additional layer of abstraction is introduced on top of the precinct data-bins by grouping sets of data-bins into blocks, much like code-blocks are grouped into precincts, and two new fields are specified, the *request* and *sblock* fields. The *sblock* field is optional and can be modified by the server if specified in the client request. The server response contains the data blocks and these are cached by normal web proxies, earning a saving in header information. The proxies can then serve other clients out of the image data already cached.

Leveraging HTTP proxies already in place allows caching of JPIP response messages with some minor changes to the JPIP client and server but without requiring any modifications to the proxies themselves. However, we wish to implement a proxy server which is JPIP-compliant and able to actively respond to JPIP request from clients. A JPIP proxy acts as a JPIP server when responding to requests from clients, but acts as a JPIP client if it is unable to completely satisfy the requests solely from its own cache. JPIP subscribes to the principle that the server knows best [1]. Thus by having a JPIP proxy which is also an active JPIP server, we can build a proxy which serves image data in a manner which best suits the client's needs from a cache which may be incomplete, while also managing the cached data intelligently.

Chapter 3

Implementation of Imaging System

3.1 System Design

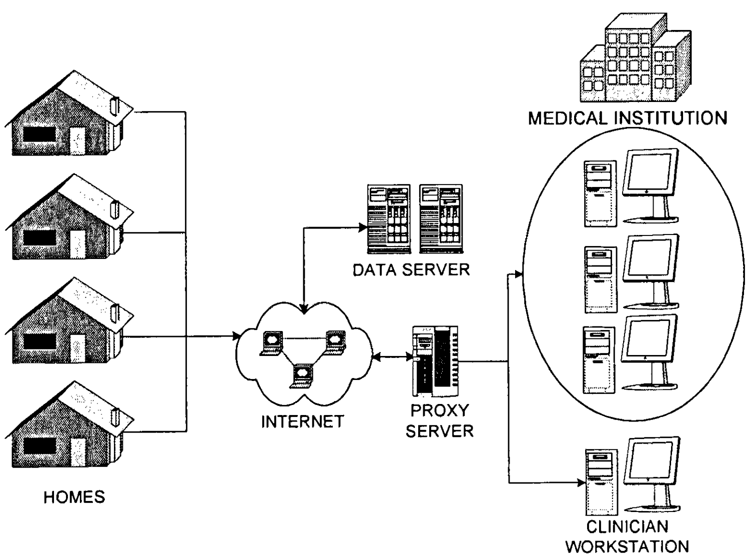

Figure 3.1: Structure of imaging system

3.1.1 Image Acquisition and PTP

In line with the aim of providing a low-cost yet effective imaging system for home telecare, we endeavour to utilize commercial off-the-shelf devices for image acquisition in our system.

Digital cameras are commonly available from many vendors, but we need to interface with the camera firmware to develop custom software for the system. To achieve this, we use an open standard called Picture Transfer Protocol (PTP) or formally known as PIMA/ISO 15740 [22]. PTP defines a set of containers and data formats, standard imaging referencing behaviour, operations, responses, events, device properties and datasets for transferring data to and from a digital still photography device and devices such as computers and printers. PTP also defines optional operations and formats, and mechanisms for extending the protocol by individual vendors. This imaging protocol removes the need for camera-specific drivers. PTP supports common device controls like image capture and image download, providing us with flexibility in application development. The PTP standard is also transport-independent and can be used over serial, USB, IEEE 1394 (FireWire), infra-red interfaces and RF interfaces like Bluetooth. The new standard provides details needed to use PTP with USB. The use of PTP is primarily advantageous because it eliminates the need to use proprietary protocols for different device vendors to communicate with the camera firmware.

As is mentioned again in Section 3.1.6, we use the open source libraries LibUSB [20] and LibPTP [21] for communication with the camera imaging device. LibUSB is written in a platform-independent manner; however LibPTP was written in the C language for use on the Linux platform. This necessitated some extra effort to port LibPTP over to the C++ language and the Windows platform.

Effort also was invested into imposing an object-oriented design on the largely procedural structure of the existing library. This object-oriented design is shown in

Figure 3.2. An abstract base class is defined as *ptp_transport*. This class represents the transport component of communications with the camera, thus in line with the specified aim of implementing PTP over any transport, this abstract base class can be inherited and its virtual methods must be implemented for each of the transports that we wish to implement. In our particular circumstance we port the pre-existing custom USB functions in the LibPTP and encapsulate them in a *ptp_usb* object which inherits from *ptp_transport* as mentioned. This *ptp_usb* object then contains the methods needed to read from and write to the imaging device over USB, as well as retrieve events and possible error messages. The core object that provides the PTP functionality is the *ptp_manager* object. This object contains the methods that implement the PTP operations as defined in the standard, for instance the 'initiate capture' and 'get object' operations. In order to use the *ptp_transport* object with *ptp_manager*, *ptp_manager* has a pointer to the abstract *ptp_transport* object as one of its member variables, and we invoke the initialize method of *ptp_manager* with a pointer to an instance of an object inherited from the abstract *ptp_transport* object (in this case the *ptp_usb* object). Then *ptp_manager* may access the transport object's methods for communicating with the imaging device using PTP. For a device which may have proprietary PTP operations as is the case with different vendors, we may inherit from *ptp_manager* to obtain a new object that adds those proprietary operations. In practice a wrapper application or object is written which instantiates both the transport object and the manager object (possibly vendor specific), initializes the manager object and then calls the *ptp_manager*'s methods, or equivalently PTP operations, as required.

All major camera vendors have provided support for PTP in most (if not all) of their products since 2002 [21]. However we are still subject to the limitation of the subset of PTP operations and properties that are supported by the camera firmware. Of the initial cameras used for testing, none supported the 'image capture' function, fortunately however all supported the 'image download' function, as well as other rudimentary operations like 'get image information', 'get image size' etc., which suf-

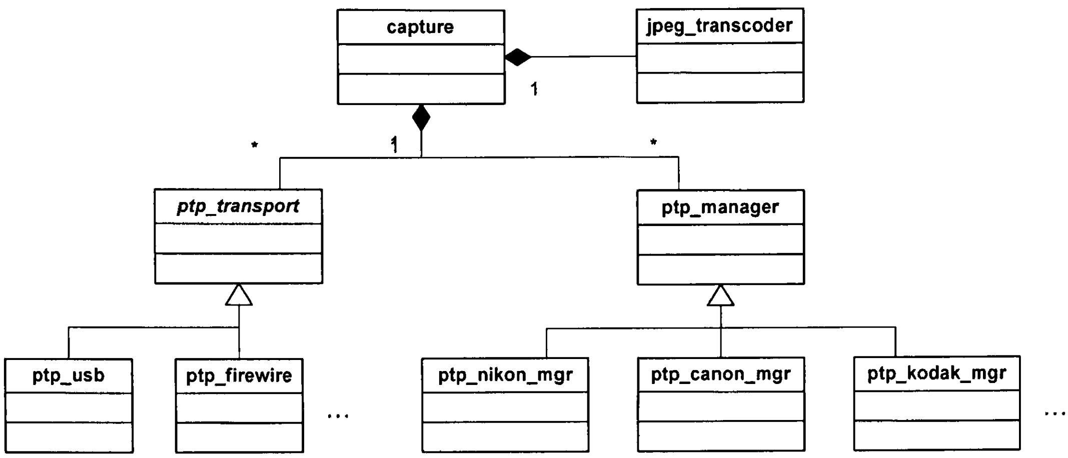

Figure 3.2: Object-oriented structure of modified PTP library

fices for our programming needs. A recent acquisition was a high-end Canon Powershot Pro1 camera, and it was found to support almost all the PTP functions and properties specified in the standard, including the important 'image capture' function. In addition, the camera supports a host of proprietary (Canon-specific) operations and properties, which provides us with even more functionality than that given by the standard PTP operations. However, although PTP is widely supported, vendor documentation about the properties and functions that a particular device supports is virtually non-existent. Frequently a device has to be queried to obtain the set of operations that it does indeed support. Furthermore, as proprietary operations and their purpose and use are by definition not included in the PTP standard, the lack of vendor documentation means that we have to rely on open source projects for information about the proprietary operations, which are also scant. This impedes the development of custom software as much time is spent deciphering the proprietary operations and properties for our own use.

As a demonstration of the industry's commitment to the PTP standard, the Camera and Imaging Products Association of Japan has released an extension of the PTP standard which covers the use of the protocol over TCP-IP networking environments, known as PTP-IP [23]. This extension enables multiple simultaneous connections to exist amongst digital imaging devices, but the most important added functionality provided by this extension is the use of PTP over wireless LAN networks, thereby enabling wireless communications to take place amongst digital imaging devices. First generation wireless LAN cameras from Kodak, Canon and Nikon have already been released as of early 2006. With regard to Telemedcare, the application of PTP for wireless devices will undoubtedly be an enhancement to the present system, and will likely be considered for implementation in future.

3.1.2 Conversion of Image Formats

The most prevalent imaging format in use with imaging devices is JPEG. With more expensive top-of-the-range devices, the RAW format (a format which may be proprietary to individual vendors, where the image is saved at a bit-depth of 24 bits for each pixel without any compression at all) is also supported, but we make the assumption that we will be working primarily with the JPEG format. In order to exploit the benefits of JPEG2000, the image capture program has to incorporate a means for converting image file formats supported by digital cameras to JPEG2000.

Commonly used web browsers like Opera, Netscape, Firefox and Internet Explorer have in-built capabilities to render and display JPEG images. During the transcoding operation, we also use the JPEG library to produce a low-resolution thumbnail of the original image captured by the imaging device. We can thus take advantage of the browser capabilities to easily present previews of the images stored in the central database to the client.

The Independent JPEG Group (IJG) [30] provides a stable and efficient software for working with JPEG images. Setting up a JPEG image for decompression with the IJG software is straightforward. To do so, a JPEG decompress object is created. This object is associated with a standard input-output source, in this case the JPEG image file. The JPEG header is then read in, and the color space is set to RGB which is the default anyway. For the thumbnail creation, the source JPEG image will have to be scaled to a suitable level, and this is achieved by setting the scale numerator and denominator. The IJG software currently permits ratios of 1:2, 1:4, 1:8 and 1:16, thus we are only allowed to scale the source image at most to a sixteenth of its original height. In practice we loop until the output height of the thumbnail is 256 pixels or less, an arbitrary figure chosen for our current implementation which seems reasonable, or the scaling limit has been reached, whichever occurs first. The JPEG compression quality is set to 50, where 100 is the best quality. These settings typically produce thumbnails of between 5 to 20kB in size.

For transcoding from the original JPEG image to the target JPEG2000 image, some coding parameter choices are explained. The target is untiled, as explained in section 2.3.2, and obviously a quality layer progression is selected (i.e. the precinct packets will be in order of increasing quality in the codestream). The number of quality layers is chosen to be between 12 and 24. This number provides a reasonable amount of layers for which there will be a perceived benefit when delivering the image progressively by quality. The number of DWT levels is chosen to be 5. This affects the number of resolutions available to the image, and practically affects the minimum size the image can be 'zoomed out' to. The precinct sizes and code-block sizes are chosen to be 64×64. This simplifies the serving policy as explained in section 4.2.2. The number of components in the target image follows that of the original image, which would be 3 (RGB). Regarding the quality layers, the lowest quality layer is chosen to have a target bit-rate of 0.01 bits per pixel, while the highest quality layer is chosen to have a target bit-rate of 2.0 bits per pixel. The layers in between have target bit-rates that are spaced logarithmically by Kakadu. This choice of bit-rates typically results in an image of about 1MB in size.

3.1.3 Storage of Images

The Telemedcare system uses a hybrid system for storage of biomedical signals acquired from the clinical workstation. In a hybrid system, data are collected and stored in 'suitcase' files in lieu of a local database at the home desktop, then periodically synchronized with the central database server through the Internet or a local area network. In this way the system is able to utilize network resources more efficiently and support multiple users [10]. The central database itself stores the clinical data, including ECG and blood pressure readings, as 'long text' in a relational database structure. This also allows more efficient organization and maintenance of the data. We leverage the database system by storing the JPEG thumbnails as well to keep the database compact, and to be able to generate previews with only regular HTML, Cold Fusion markup, and SQL expressions. The use of JPEG thumbnails allows the preview of many images at once with only a small cost in bandwidth. There is thus no need to send the full JPEG2000 image if it is not required.

In order to keep the existing database small, the actual images are stored separately from the database. We envision changes to the way the actual images are accessed and manipulated; these changes would be made simpler if the actual images are separate from the existing database.

The 'central database server' is actually an umbrella term for a collection of servers: an object broker server, a processing server, and a database server [10]. The object broker serves as a directory service that controls the load and failsafe properties of the system. The object broker server can redirect requests to other processing servers in the event of the unavailability or overloading of a particular processing server, thus improving the response of the system. After being directed to a suitable processing server by the object broker server, the client sends data and commands to the processing server. The processing server now establishes a connection to the database server, the main data storage component out of the three types of servers available. The processing server is so named because data processing and analysis is performed here. The processing and analysis converts the raw data to information to be presented to the end client.

The processing server communicates and connects with the database server through the Microsoft ActiveX Data Object (ADO) application programming interface. ADO is a Component Object Model (COM) object which acts as a wrapper for Object Linking and Embedding (OLE) DB and Open Database Connectivity (ODBC) data sources, providing programming languages with a higher level of methods and properties for accessing these data sources. COM is another Microsoft technology which enables cross-application communication and dynamic object creation in any programming language which supports this technology. OLE DB and ODBC in turn allow access to data sources (which may or may not be databases, for instance spreadsheets), in a standard way. ADO thus allows customized programs to access data from different database types. Although ADO can be used to issue arbitrary SQL commands, this is not necessary when using ADO. ADO consists of several top-level objects, but the most important are the connection, recordset and command objects. A connection to the data source is established via the connection object, and data can be retrieved and stored in the recordset object. The command object is used to issue database commands or SQL to the data source.

ColdFusion [24] is a web application development system based on the tag-like ColdFusion Markup Language and scripting. In the Telemedcare system, the ColdFusion server resides on the processing server component of the central database server. It is primarily used in this circumstance as a database querying and management system, as well as a session, client and application management system. A dedicated web server is tightly coupled with the ColdFusion server and together they are set up to coordinate client request handling.

In the present implementation, the image data are inserted into the database server as with other medical data that is generated by the Home Desktop. After synchronization the processing server will open a connection to the database server, process the image records and store the main JPEG2000 images in a separate directory on the processing server, while the thumbnails will be reinserted into a separate table as part of the database server. Other metadata related to the images can be inserted together with the thumbnails as part of a comprehensive collection of information related to the patient images generated. These will include a link to the path of their JPEG2000 counterparts in the processing server, which will be served separately. Now, the placement of the main JPEG2000 images in the processing server is strictly temporary. We envision a dedicated image server running the Kakadu image server as a long term solution, in order not to impinge on the processing server's role of converting raw data into information to be presented to the end client, and burden it with an additional task of housing and serving the JPEG2000 images. In fact, we envision a completely separate image retrieval system to go hand-in-hand with the image server as well. However, for the moment we piggyback on the existing relational database in use to store retrieval information about the images, and on the processing server for actual storage of the main images.

3.1.4 Image Database Design

The image thumbnail database has a simple relational design (Figure 3.3). The database consists of four tables (the patient table with the table ID exists previously). The table 'Image Info' contains a primary key 'Image ID' and a foreign key 'Patient ID', together with other relevant original image information such as the image height, width, and original size. The table 'Image Thumbnails' contains the actual thumbnails stored as BLOBs. The table 'Image Modifications' stores the history of modifications that an image has undergone. The type of modification is stored in the 'Modification' table. Finally, the location or directory path of the actual non-thumbnail image is

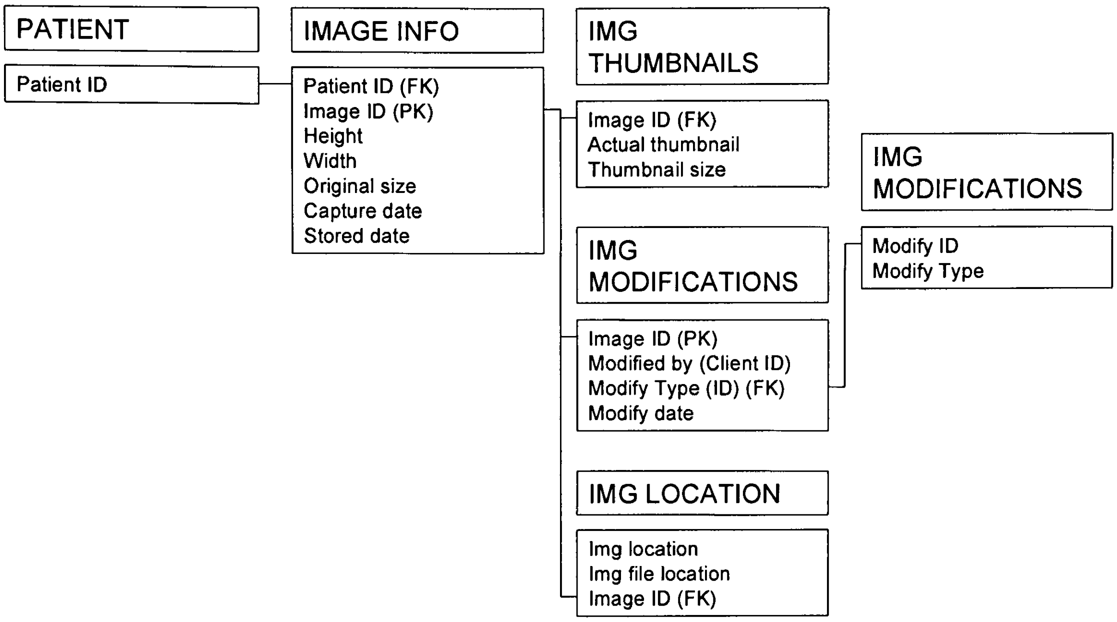

Figure 3.3: Image Thumbnail Design

stored in the 'Image Location' table. The primary keys 'Image ID' and 'Modify ID' are used to link the tables together to form the necessary relations.

We note at this point that this database design is yet to be fully implemented due to time constraints.

3.1.5 Image Metadata

From Section 2.2.2, we know that we are able to insert metadata into a JPEG2000 file in the form of XML, where the elements are well defined in the standard as described in a previous section. Kakadu (see Section 3.1.6) provides functions for us to insert and extract XML metadata easily, as XML is nothing more than mere textual data. This provides us with tremendous convenience in metadata handling, not only in terms of storage but also in the transmission and delivery of the metadata. By embedding the metadata in a JPEG2000 file within an XML box, we ensure that the metadata is tightly coupled to the image itself, and indeed the contributors to the

JPEG2000 standard most probably had this concept in mind when designing the image file format. When we deliver the image file from the Home Desktop to the central database, the metadata are simultaneously transported and stored.

When prompted by a client request, the metadata stored with the image file are extracted to a temporary file. But how do we process the XML that has already been extracted? We utilize the power of ColdFusion to parse the XML and transform it with an XML stylesheet to a form more suitable for presentation [34]. This is easily accomplished by a few function calls which encapsulate most of the XML processing capabilities of ColdFusion. The first of these is XMLParse, which accepts a variable containing raw XML text that has been read in previously, and parses the elements and attributes of the XML text internally. The second of these is XMLTransform. This function takes in the variable containing the parsed XML content, as well as a variable containing the XML stylesheet that we wish to use to transform the XML content. The function then does all the hard work, performing the necessary transformation to produce a HTML file that will be returned to the client.

We follow the XML Schema strictly for the metadata elements that are to be inserted in the image file. These will mostly be from the image creation group of metadata, which records the camera settings that were used in the acquisition of the image. Other metadata groups defined in the standard, i.e. metadata history and content description metadata will not be included in the current implementation, as these metadata are usually populated through user input.

3.1.6 Software and Hardware Requirements

To work with JPEG2000, we use the Kakadu Software [6]. Kakadu is a comprehensive, optimized and fully compliant software for developing JPEG2000 applications. Kakadu is not freely available, but a licence can be obtained at minimal cost, and the licence grants the user with unlimited upgrades and compilations.

From [6], Kakadu is a complete implementation of the first part of the JPEG2000 standard, Part 1. Kakadu also provides a comprehensive implementation for several of the most useful features from Part 2 of the JPEG2000 standard. For our purposes, Kakadu provides the following useful features:

- Kakadu is able to perform compression and decompression of image files in various formats, from small images to images in the tera-pixel range. Kakadu can decompress and render almost any JPEG2000 source to memory or to an output image file.
- Kakadu can transcode between related representations, performing operations which are natural in the context of JPEG2000.
- Kakadu can decompress and render images interactively.
- Kakadu provides support for interactive client-server applications by implement part 9 of the JPEG2000 standard involving JPIP.

Development is being performed on a Windows XP system. We code in C++, on a MS Visual C++ 5 environment. Each of the proxy, patient and client applications uses or is based on Kakadu. For the patient-side camera capture utility, we use the open source LibUSB and LibPTP libraries and a LibUSB wrapper to allow custom USB communication. If the platform is non-Windows XP an additional open-source generic digital camera driver must also be installed in order to talk to the imaging device. The LibUSB and LibPTP libraries are available via [20] and [21]. We use ColdFusion markup and SQL for database queries.

The applications themselves have been executed on Windows 2000 and XP platforms. A precompiled Kakadu DLL is required to run the programs. The home desktop modules are programmed in Borland Delphi 7, making use of visual components (buttons, forms, labels etc.) that have already been created for Telemedcare. The processing server applications for the imaging system are also created in Delphi, this time utilizing the ADO objects provided by Delphi.

There are no special hardware requirements for running the applications, except for a personal computer with internet access (dial-up, broadband) which already exists as the Home Desktop, as well as a digital camera for imaging. The only caveat is that the camera needs to support PTP.

3.2 System Operation

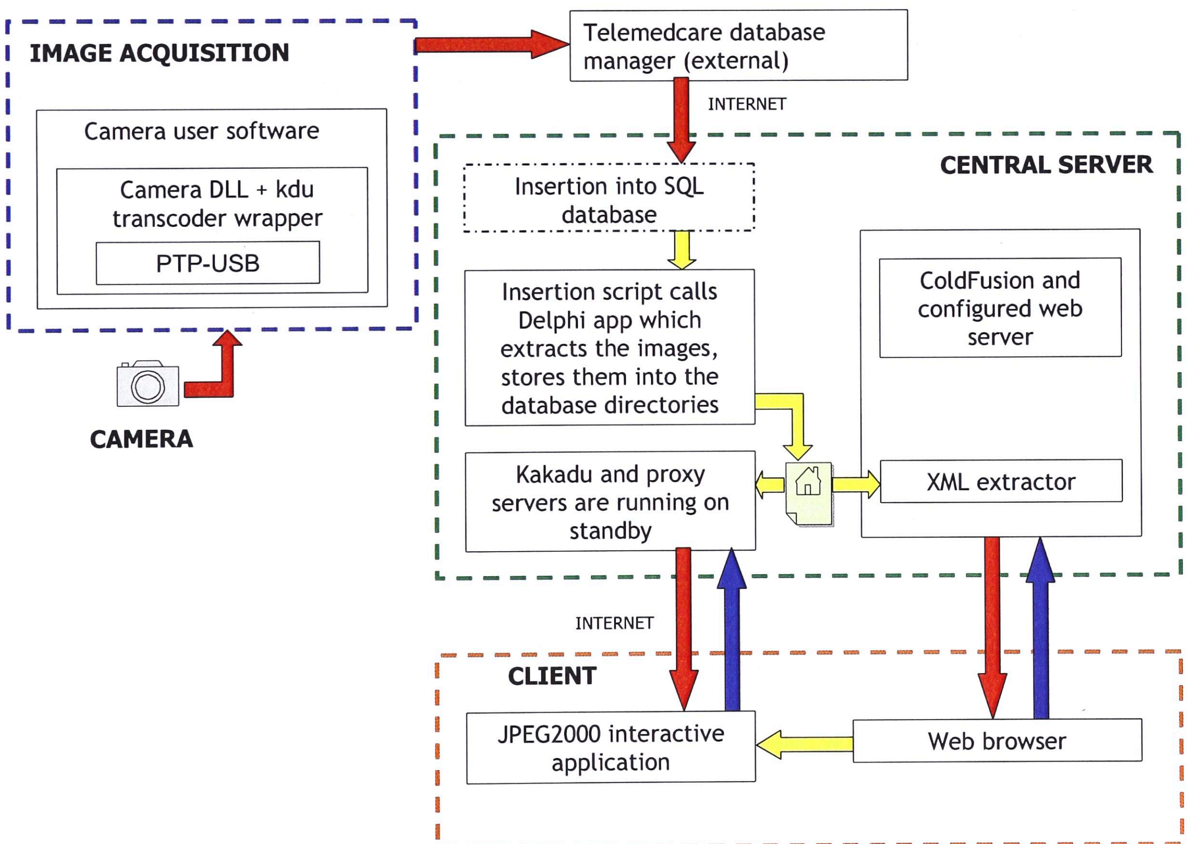

Figure 3.4: Interaction of system software components

3.2.1 Home Desktop

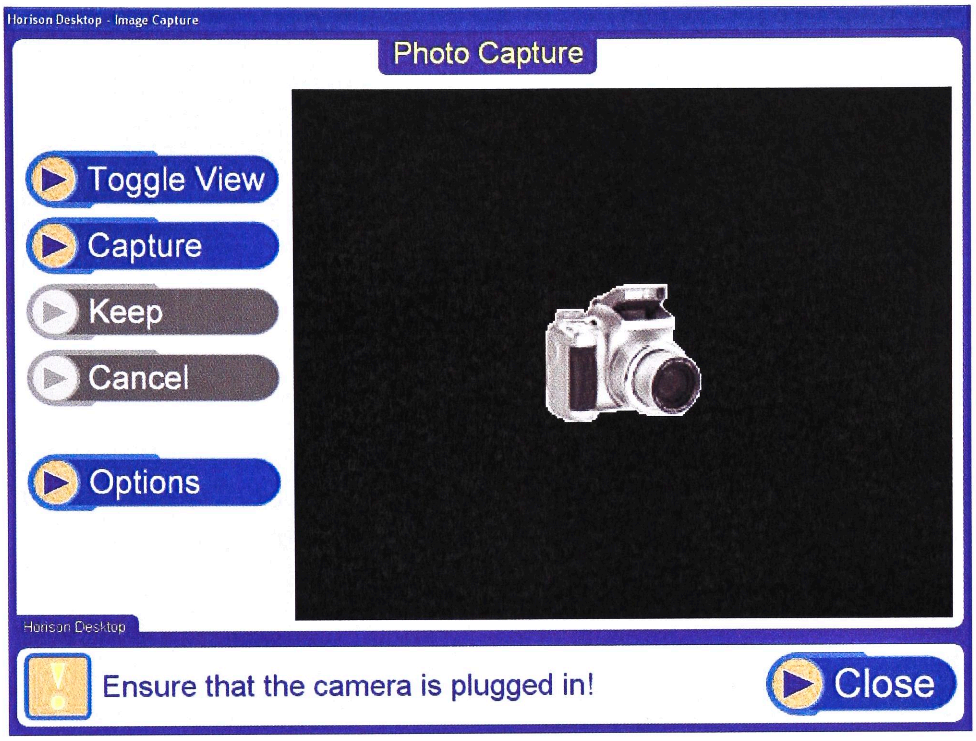

Figure 3.5: AcquirePhoto (ready state)

At the patient side, the home clinical workstation presents a simple user interface to the patient. The patient captures an image using a digital camera connected to the workstation. At present the transport used is USB. The image is captured as JPEG but converted into JPEG2000. A JPEG thumbnail is created simultaneously. The image is then encoded in a Multipurpose Internet Mail Extensions (MIME) encoding and sent as an XML file to the Telemedcare database, along with the thumbnail. The use of MIME and XML are strictly to conform to the existing Telemedcare paradigm of data encapsulation on the home clinical workstation side. Note that this usage of XML is separate from previous discussions about XML usage concerning image metadata.

For the image acquisition, an image is acquired by the AcquirePhoto module as part of the Home Desktop system. Figures 3.5 and 3.6 show the image module user interface, and Figure 3.8 shows the block diagram of the module program.

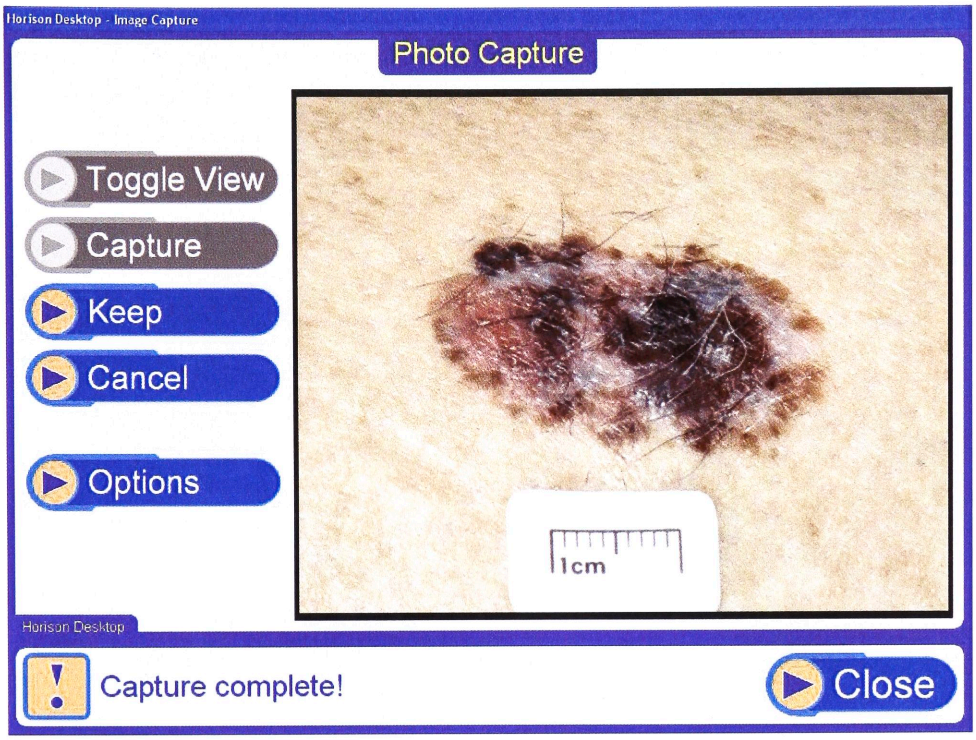

Figure 3.6: AcquirePhoto (after image capture)

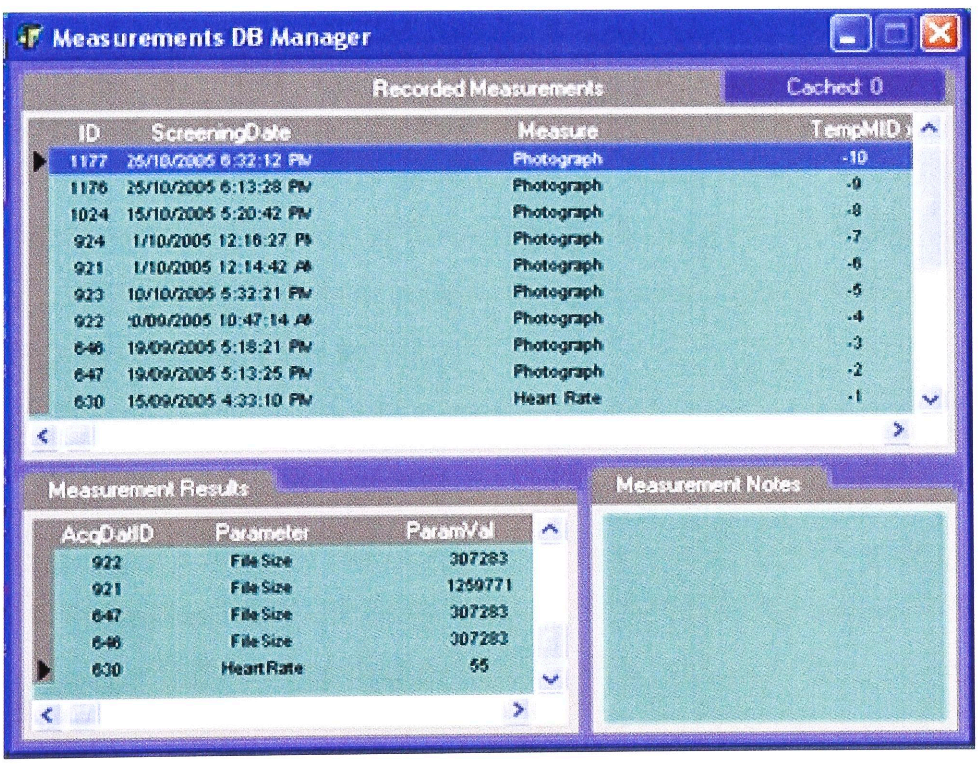

Figure 3.7: DBMeasureManager module

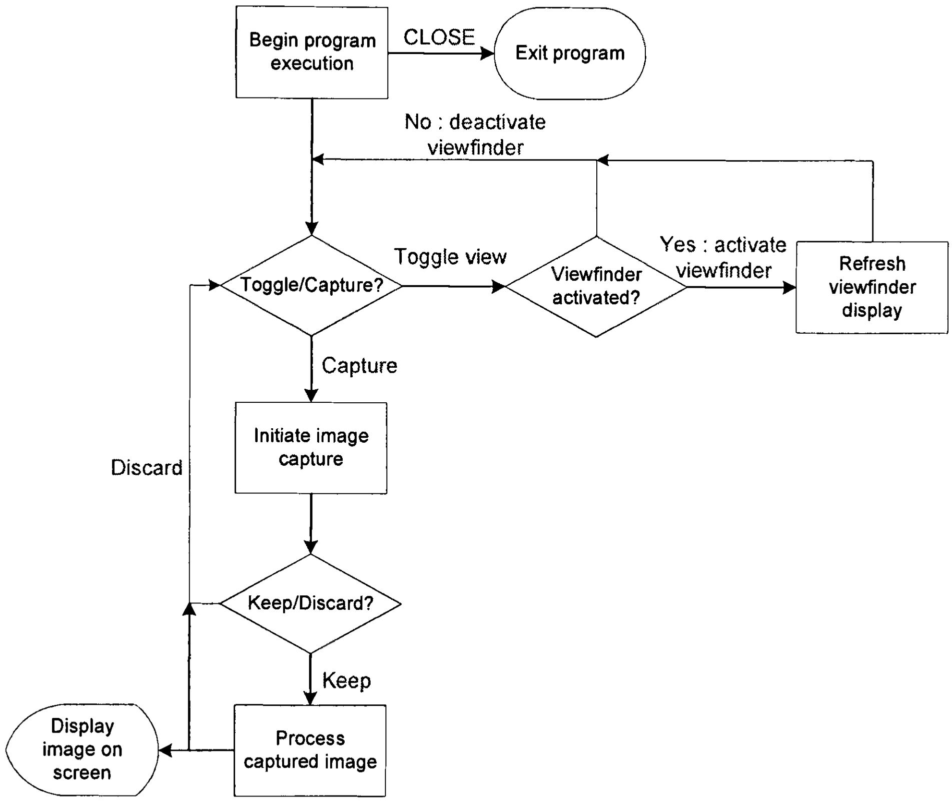

Figure 3.8: AcquirePhoto module program structure

Upon start-up, the user is able to capture an image, or to toggle the viewfinder mode. When the viewfinder mode is turned on, the screen shows a continual display of images as seen from the lens of the camera, much like a live video display. When the user is satisfied that conditions are suitable for an image capture, the capture button may be pressed to commence capture.

When the patient clicks on the 'capture' button, the software performs an initialization step. This involves looking for an imaging device at the workstation's USB endpoints, and if a device is indeed found, setting up a custom PTP structure which contains the read and write functions to be used (this makes the structure transport independent) as well as configuring the device's USB interface. This structure will be used in all PTP transactions involved in the capture process. A PTP session is then opened with the device. A session is needed for PTP transactions, except when we wish to obtain the device's properties. The captured image has to be stored locally on the camera, so the storage media IDs are obtained from the device by asking the device to populate a PTP storage ID container.

The device now carries out the capture transaction, and waits upon its completion. Specifically, the program waits for the device to return a PTP event container with the capture image object ID. The program uses this ID in two ways. Firstly, it requests the device to populate a PTP object info container, in order to obtain the size of the compressed image. This size info will be used in the image transcoding operation. Secondly, it requests the device to return the image itself in a memory buffer. One unavoidable aspect of downloading the image from the device is that it does so in its entirety, thus the program has to be prepared to allocate a large amount of memory during the capture process, as much as the compressed image size. (However, some newer cameras such as those from Canon do support a new PTP operation which allows one to get an object partially, thus potentially saving memory if the operation was repeated until the object was completely obtained. This was not implemented in time.) Once the capture process is successful, the PTP session is closed and the USB device is released.

In the transcoding operation, we set up a JPEG decompress object as the source and a JPEG2000 compress object as the target. The JPEG image obtained from the device is actually written to a file first, as it is more convenient to associate the file with the JPEG decompress object. At this point, the memory buffer containing the original JPEG image can be recycled if desired. From this point the memory structures used during transcoding will be as large as one raster scan line of image, multiplied by the number of components in the image (typically three for RGB). For the JPEG2000 compress object, we use the following parameters: rates of 0.01-2.0 bits per sample, 20 quality layers, 5 resolution levels, precinct and code-block sizes of 64×64 (i.e. each precinct contains exactly one code-block), and using a reversible transform. The decompression and compression operations happen one after another and incrementally, i.e. when a row of pixels are decompressed by the JPEG object, it is pushed into the JPEG2000 object. The JPEG thumbnail is similar to the transcoding operation, except that the compress object is another JPEG object. The original JPEG image is then scaled and written out to the compress object.

From here on at the patient side, the existing Home Desktop software infrastructure is used to perform synchronization with the central database server and transmit the image data over the network.

The main image and thumbnail data are MIME-encoded in Base64 format in preparation for network transmission. Briefly, Base64 involves taking every 6 bits of the original binary image data and mapping the 6 bits to an 8-bit character code, in effect expanding the original data by about 33%. The module then packs the image and thumbnail together into XML format, but they are encased in different tags. This image in XML format is sent to another module of the Home Desktop, the so-called 'database manager' module, which takes responsibility of accepting data from all other modules of the Home Desktop and synchronizing them with the central database server. The data are sent to the database manager module by way of Windows message passing function calls.

The database manager module then accepts, compresses and packs the data in a so-called 'suitcase' file. The compression utility that is used by the module uses the Burrows-Wheeler Transform as the heart of its compression technique, which gives superior compression rates compared with commonly available utilities such as Winzip and WinRar, which typically use dictionary methods of compression. Once the XML data are compressed and packed in a suitcase file with an extension of '.sut', the file registers on the list of files ready to be sent by the database manager module. When the module is called to synchronize the files in its cache with the central database, the module opens a connection with the central database server and duly sends the files to the server. We wish to note here that passing the data from the image acquisition module to the database manager module takes an unduly long time. By examining the reverse process at the central server (i.e. the extraction of the compressed XML record to obtain the original image data), the delay can be attributed to the use of a DOM implementation object by the database manager module to build the XML file with all the requisite tags. The delay is unique with regard to image data as they are much larger than the data generated by other modules, for instance from the ECG module. At present, this delay is deemed acceptable, the rationale being that a patient is likely to generate only a few photos at a time, and that the workstation used by a patient is entirely dedicated to the Home Desktop, meaning that the patient does not use the workstation for other purposes. This issue will be revisited in the next section.

3.2.2 Database

During episodic synchronization with the patient workstations, the central database server accepts incoming suitcase files containing the biomedical patient data and images, together with the thumbnail previews, and inserts them into a table called 'AcquiredData'. All newly-accepted data files will be stored as records in this table, including our image data. The record which appears in the table remains the compressed Base64 data encased in XML format. So, we need to reverse all the processes which were involved in preparing the image data for transmission on the Home Desktop side. The processing server runs a Cold Fusion server which parses the web requests from clients, and accesses the database server accordingly. The actual image requested is served by a JPIP server running separately. The client application establishes a stateful JPIP session with the server.

After synchronization, the processing server runs an executable to process all image records, including the newly inserted data, since the last synchronization. Firstly, the data are retrieved from the database AcquiredData table. Now, ColdFusion has the capability of performing SQL queries itself. However, retrieving a record from the database and storing it as a ColdFusion variable does not work simply because the Data field in the AcquiredData table (which is the field containing our image data) is a defined as a field of Long Text. Therefore if ColdFusion is used to retrieve the record, it recognizes the data as text and duly adheres to its default character encoding set, which means that characters which are not valid in that character set gets replaced with the 3F byte (the '?' symbol, appropriately). Obviously this means that the data will be corrupted, even if one byte is replaced in this way, and there are no 8-bit character sets for which there is a valid character for each of the 256 representations. If a representation is invalid it usually is a control byte of some sort. So this rules ColdFusion out, and we have to resort to using other means of retrieving the record from the database (which ought to be an extremely simple process).

We do this by means of ActiveX Data Objects and methods. Borland Delphi implements the ActiveX Data Object library which allows us to use ADO methods to establish a connection with a database and retrieve data using SQL. The character encoding problem which was encountered with ColdFusion does not apply here, and the image data are retrieved in pure binary form and written out to a temporary file. The record is then removed from the AcquiredData table.

The executable then continues on to process the temporary file containing the image data. The application decompresses the compressed data in the file using the same library as that used by the database manager module on the Home Desktop. The data has now reverted back to XML form, so the XML tags are now parsed to obtain both the main image data and thumbnail image data. These are both written out as temporary files again, while the temporary file containing the compressed data is deleted. The script now decodes the Base64 encoding of the two temporary image files, and we are left with the original image data for both the main image and thumbnail. The main image is written out to a special directory on the server and the temporary files are deleted. Once these processes are complete, the images are ready to be served to remote clients.

File size (kB)	Time (s)
19285	<1
307283	57
1259771	1088

Table 3.1: Processing time for preparation of image data for transmission for different file sizes

File size (kB)	Time (s)
19285	0.24
307283	0.90
1259771	3.11

Table 3.2: Processing time for preparation of image data for transmission for different file sizes, using custom function

Some simple tests show that as the size of the image grows, the server-side processing time grows in an exponential fashion. This would also apply for the processing time of the Home Desktop to prepare the image data for transmission as it is essentially the reverse of the server-side processing, as mentioned in the previous section. The following statistics were obtained:

The time to decompress the largest file used above (1.26MB) was about 1 second. The ADO retrieval of the record and the Base64 decoding also have negligible processing time. Thus, the delay in processing can be attributed directly to the object which parses the XML. This was a concern, as in a real situation the delay of sending the image data from the Home Desktop to the central database server would be further compounded by the network that is used, more so if the network is one of low datarate like dial-up. It must be noted however that synchronization between the central database server and the Home Desktop usually happens at night, when time is not as crucial. Nevertheless, the use of the DOM implementation object which was used to parse the XML was discontinued, and instead, noting that an XML file is essentially a text file, a custom procedure was written to extract the image data from the correct XML tags. Some new results were obtained:

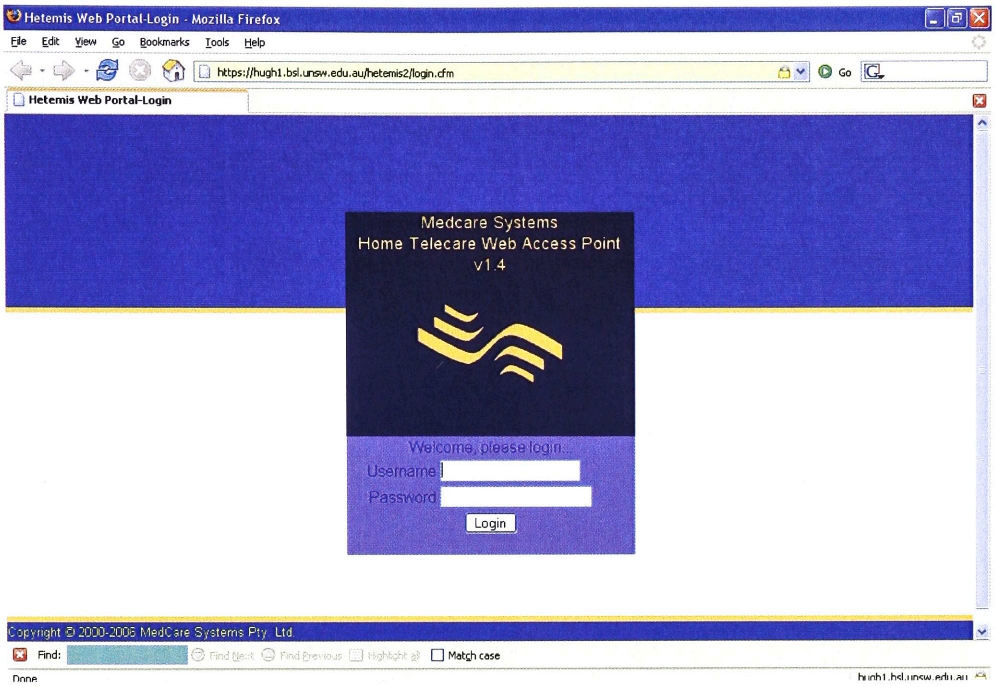

Figure 3.9: Telemedcare client web interface 1

These processing times when using the custom procedure are a drastic reduction of the previous processing times using the DOM implementation object!

3.2.3 Client-side

At the client or GP workstation, the client browses images available on the central database through a web interface, which is already in place for the existing health measures available (Figures 3.9-3.11). The client is able to access the web interface through being provided a username and password. After selecting the patient whose images the client wishes to view, the database returns the thumbnail previews as requested. When the client clicks on the preview, an application is launched which allows interactive browsing of the requested image via JPIP. The client may be served by a proxy cache, but this would be transparent to the client.

As mentioned in Section 3.1, the thumbnail previews are generated using regular

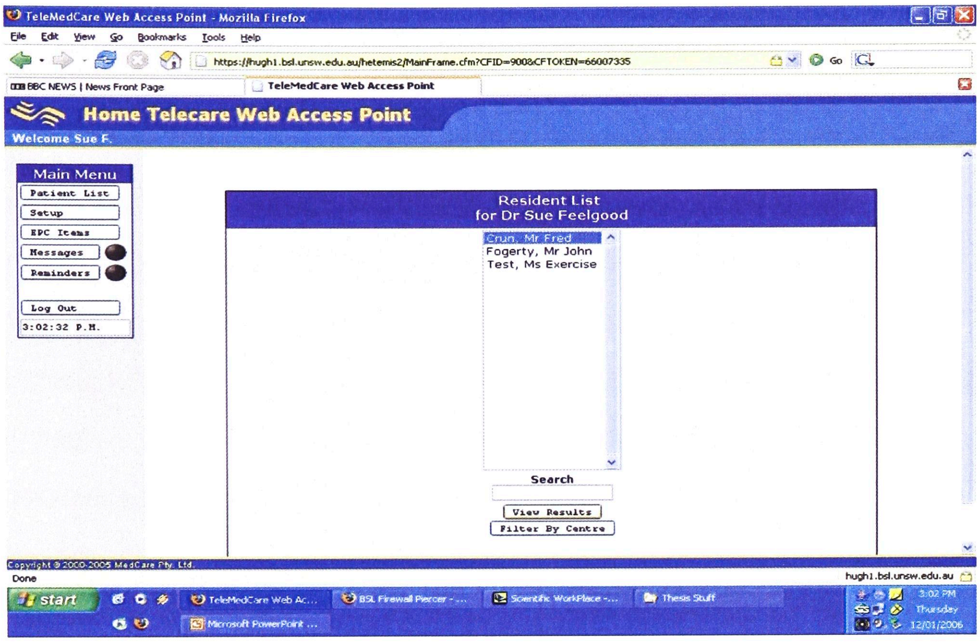

Figure 3.10: Telemedcare client web interface 2

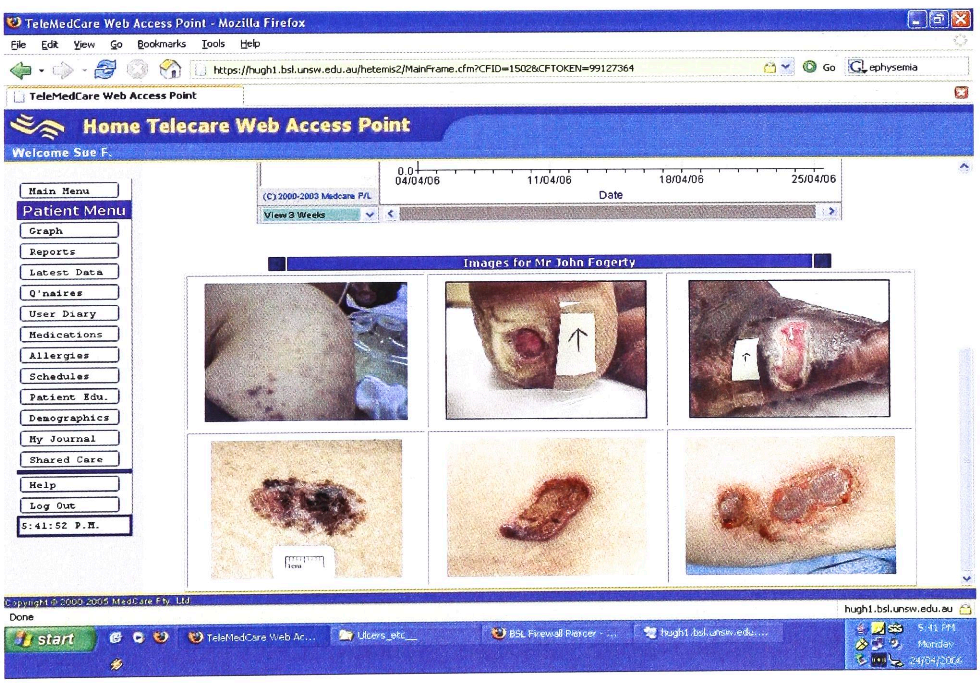

Figure 3.11: Telemedcare client web interface 3

HTML, Cold Fusion markup and SQL queries. The thumbnails fetched from the database are written to a temporary location for the web server to deliver to the client.

The preview images are accompanied by some auxiliary information about the image, for instance the image name, file size, image dimensions, date captured and date stored. Other auxiliary information can be trivially added to the dynamically generated preview page if necessary. One particular type of information which will be of interest could be a brief description by a GP of the image and any clinical symptoms it may depict. Another could be the revisions that the image has gone through, e.g. any annotations or image processing performed on the image.

At present, to demonstrate the delivery and presentation of metadata, a simple hyperlink is added below the thumbnail image on display and the JPIP link to the image. When the hyperlink is clicked, a request is sent to the server to extract the metadata from the XML box(es) of the desired image. The request is handled by a ColdFusion script which first launches a helper application which opens the image file, searches for the XML boxes, and writes the contents to a temporary text file. The script then reads in the XML data from the text file and parses the XML elements with just a single function call. Additionally, the script reads in the XML stylesheet which will be used to transform the metadata for a form more suited for presentation, i.e.HTML. This is accomplished once again with a single function call, utilizing the power of ColdFusion. The transformed XML is served to the client in a separate pop-up window.

3.2.4 Interactive application

The interactive application is built on the *kdu_show* application provided with Kakadu to view JPEG2000 images. Besides being able to display local or remote images within a selected region-of-interest if specified, *kdu_show* is also able to edit optional metadata such as labels in the images, and perform simple spatial manipulations like rotate and flip. The windowing structure of *kdu_show* was modified to have the main display on the right with a black background, and some other windows on the left which purpose as yet is undetermined (could be some thumbnail display of local images). We plan to add two major forms of functionality to the application: simple image processing tools and the ability to annotate the images displayed using the application.

3.2.5 Proxy cache

Clients are served by a proxy server (see chapter 4) instead of the central database, which behaves as the ultimate server for all the images collected from the patients. The proxy is JPIP-compliant, i.e. it accepts JPIP requests from a client and issues its own JPIP requests to the ultimate server. The proxy caches images which it requests, which it can then use to satisfy subsequent requests for the same images, possibly from different clients. The proxy can cache and serve incomplete images. We adhere to a JPIP principle that the server may serve image data in any order it pleases. The proxy can just serve whatever is present in the image cache, or it can follow a policy which selects image packets to be served which will optimally reduce the distortion perceived by the client. Policy choice depends on the network conditions which are monitored by the proxy. The proxy is deficient in the following senses:

- the proxy is filled with image data only when prompted by client request;
- it only supports HTTP transport although JPIP provides for HTTP-TCP as well;
- and it has only been tested with meta information necessary for opening and displaying the JP2 file format.

The proxy is explained in more detail in the next chapter.

Chapter 4

Implementation of JPIP Proxy

4.1 Proxy description

This section is concerned with the development of a proxy cache server which is compliant with the JPEG2000 interactive protocol, JPIP [1], [5], [9].The JPIP proxy server developed in this project will normally be placed between the ultimate clients, usually a general practitioner or some other medical facility, and the central Telemedcare database server. The proxy server can coexist with the central server or can be placed closer to the clients, in which case the data link between the proxy server and the clients is expected to have a higher bandwidth than that between the proxy server and the central server.

4.1.1 Proxy server structure

Figure 4.2 shows the structure of the proxy server. The proxy is built on top of the JPIP server implementation in [6].

When a client request arrives at the proxy, the proxy's connection manager creates

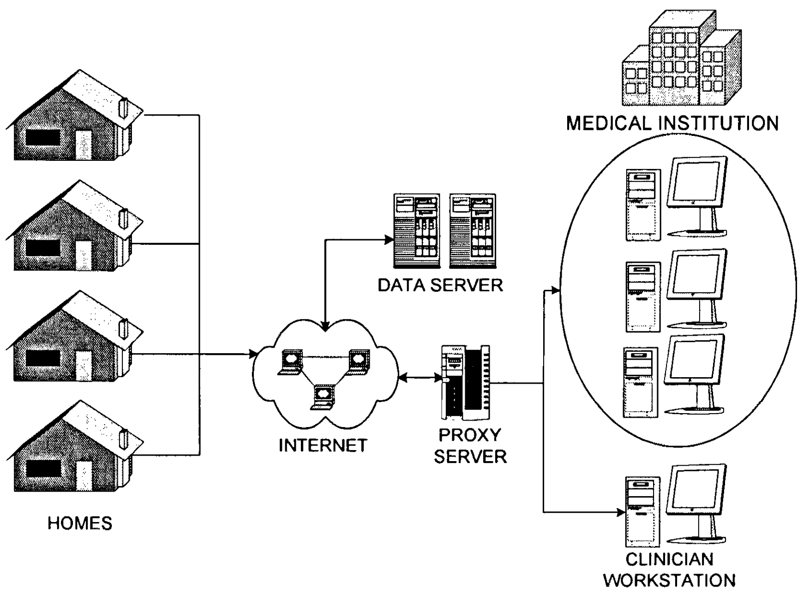

Figure 4.1: Placement of proxy server in imaging system

a new thread to service the request. In the new thread, an HTTP channel is created and a channel-ID is assigned if the client wishes to reuse the channel in a stateful session. JPIP provides additional details on using JPIP on a HTTP-TCP channel, but we only consider JPIP on HTTP here. The image source requested is opened, if it is not already being served to another client. Since many clients may access the same image resource simultaneously, the source actually maintains a separate list of thread objects which have access to it, for each client. If the session is stateful, each thread object maintains state information on the client cache contents assuming that each server response is cached by the client, but the client may issue cache-model directives of its own (see chapter 2). Note that we wish to make the distinction here between the *server cache-model*, which is state information modelling the client cache contents, and the *main cache*, which is the actual store of image data at the proxy. So, client directives affect the server cache-model, which exists as long as there is a session between the proxy and the client.

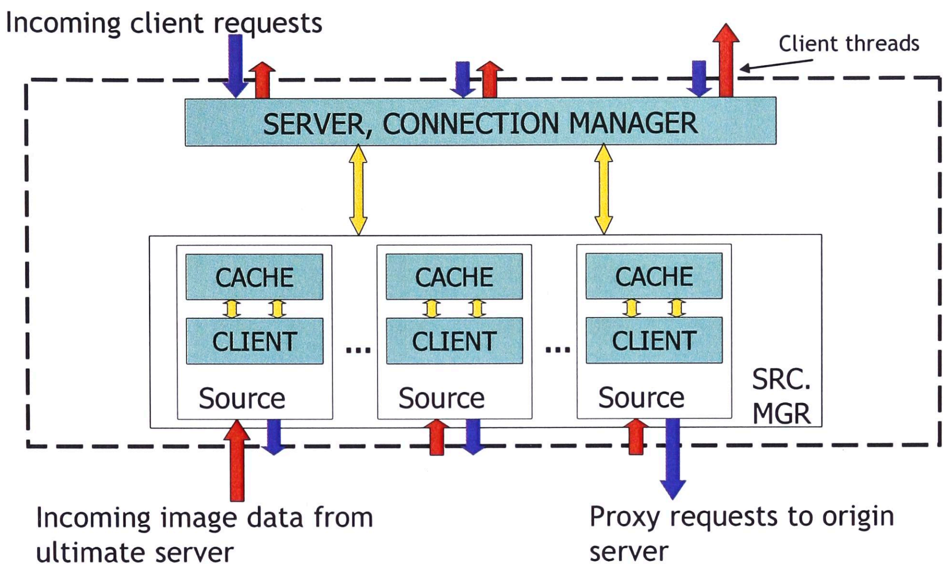

Figure 4.2: Architecture of JPIP proxy server

4.1.2 Cache data management

Image data which were previously cached by the proxy is written out to a file, when all clients accessing the image disconnect from the proxy. These data are reloaded into memory when a source is to be reaccessed by a client. The cache file and cache structure in memory are faithful to the data-bin concept. Data are written to and read from memory by treating each data-bin as a unit; furthermore, the cache is searchable according to data-bin, and we grow the cache in groups of data-bins at a time. Anytime a source is opened for the first time, we create the cache object and load the image data into the cache. This object also acts as a client to request image data from the origin server if necessary. The proxy parses the client request fields as a normal JPIP server would. Based on the information obtained from the fields, the proxy determines whether the image window specified in the request can be satisfied by the cache.

4.1.3 Serving client requests

Suppose that the image requested is not present in the proxy. In this case, the source and its associated cache object are still created, but the cache remains unpopulated. The cache object creates its own client thread and attempts to connect to the image's origin server as specified in the *Host:* header of the ultimate client's request. The proxy constantly informs the cache object of the client window requests to form the basis of its own requests to the origin server, so the requests made by the proxy to the origin server mirror the needs of the ultimate client. However, if the proxy fails to connect to the origin server after a preset time-out, or if there was a problem accessing the image (for instance the image is not present in the origin server as well), the proxy reports the failure to the client and the session is aborted. Otherwise, the connection is successful and the image is returned.

Suppose now that the client window request has no intersection with the proxy's cache, or that the image request has never been cached. We assume that a connection to the origin server has been established successfully. We also assume that an image encapsulated as a JP2 file is being requested; we serve raw codestreams in a similar way to JP2 files. We note at present that the effect of JPX files on the behaviour of the proxy is not well understood. The proxy now accepts image data contained in the origin server's responses to its requests.

When attempting to serve an incomplete image, we face the problem of empty packets which slightly complicates the serving operation. An empty packet may be generated for two reasons [2]:

- The precinct code-blocks contribute no samples to the packet;
- The intersection of the precinct and the client window has no samples (literally empty).

The proxy uses the extended header option *jpp-stream,p-type:extend* [9] to obtain information about completed packets for each databin from the origin server. This information is returned in each JPIP message header in an auxiliary variable byte-aligned segment (VBAS). This information therefore means that the origin server does the work of determining whether a packet or quality layer is complete, or whether a packet is empty. Information about complete packets provides almost the same flexibility to implement custom serving policies as is available when the image is complete, since we can at least manipulate the quality layers in each data-bin. An example is the policy of rate-distortion optimization to improve the client's perceived quality [3], [5] used in a stand-alone JPIP server implementation. We discuss more about server policies in section 4.2. However, here we use packet information to implement simple packet thresholding for every data-bin, and adopt a policy of serving any and all data which lie below the threshold. We refrain from adopting more sophisticated policies at present; such policies may likely entail a degree of waiting for data from the origin server which may be unacceptable and reduce the responsiveness of the proxy. So, each time a data-bin caches a packet and its total packets exceed the threshold, the proxy increments the threshold. Image data are read directly from the data-bins in the form of byte-ranges, packed into data chunks and delivered to the client.

4.1.4 Serving meta-data

The above discussion only applies to data-bins which actually hold image data, precinct data-bins. Meta data-bins and the main header data-bin present a different problem to the proxy. Consistent with the 'server-knows-best' principle, it is conceivable that an origin server may elect to send required meta-data for an image at any time during a session. For our proxy, we simply choose to block on the required meta-data and main header data-bins. This meta-data is contained in meta data-bin 0, or alternatively the root meta data-bin [1]. With the use of placeholders however, it is possible that some required meta-data resides in other meta data-bins. We do this because by definition the client would be unable to decode image data sent before the required meta-data. A sensible JPIP server should send this meta-data as soon as possible to enable a meaningful interactive browsing session to take place. Examples of required metadata include the JP2 signature box and the file type box.

It is not uncommon that metadata other than that required to render and display the image is interleaved with the imagery data. We note that the effect of metadata arriving in such a fashion on the proxy server has not been fully understood. This means that the proxy in its present form may not necessarily maintain the metadata structure of the image requested that was determined by the origin server. Since the proxy blocks on the required metadata to form a tree-structured record of metadata for serving, the proxy essentially ignores all these non-mandatory metadata arriving later, and we assume that the requested image only contains meta-data necessary for the correct decoding and rendering of the image. Another unresolved issue is the question of placeholder allocation by the proxy server, where as mentioned in section 2.4, a JPIP server is free to determine placeholder structure when serving an image. We decide not to alter the placeholder structure imposed by the origin server in the current implementation.

4.1.5 Sources and source clients

Currently an object manages image sources and threads associated with each source. When the server receives a request, it services the request by having the manager object find the source amongst the list of sources the manager object presently manages. If the source cannot be found, the server attempts to open the source or create a new source object and add it to the list. A new thread is created and attached to the source's list of threads that presently access the source. A source should ideally have one source object associated with it. All source threads will share the source client just as they share the source itself. It thus makes sense for the source manager to manage the source client object associated with the source. Currently we assume that the source client will be driven or activated by an incoming request for the particular source object. In contrast the ultimate role that is desired of the source client is that of a client which works autonomously, establishing a connection with the origin image source if the source was incomplete, and determining the data-bins that are required to complete the image and requesting them.

4.1.6 Assignment of target IDs and cache files

The source client formulates requests initially based on the image desired. The request is sent to the proxy server by HTTP. The proxy receives the request and searches for the file in the cache. If the file is not present in the cache, the proxy formulates a new request for the file from the original server, based on the IP provided in the client request. The original server looks for the file in its storage. If the file is not present, it sends an error message which is relayed by proxy to the client. If the file is present, a new target ID is assigned to the target file. The image is transmitted to the proxy as per the proxy requests. The proxy caches the image file as it arrives from the original server. The new cache file is assigned the same target ID as that given in the server replies. The proxy relays the JPIP messages sent by the original server in response to the proxy requests, to the client.

So, the target ID is always driven by the original server, if the file is served from scratch. In summary, the target ID is provided by the server (either the original server or the proxy server) if client is requesting an image.

Suppose that if a file that is requested by a client is not present at the cache, the image file is created, but merely as a placeholder to the cache filename. Currently the placeholder file contents consist of a string stating that it is indeed a placeholder file, followed by a newline and then the name of the actual cache filename.Then, the image is updated at the cache. The generation of the target-ID is currently based on the image filename and its last modified date. So a consistent target ID is provided as the modified date for the file remains unchanged. JPIP requests which specify a byte-range from the image file should cause a unique target-ID to be issued for each unique byte-range, and this is reflected in the inclusion of the byte-range in the placeholder filename.

Finally, we have decided to retain the structure of the data written into the cache file which mirrors the way data-bins are managed in memory, even when all databins have been fetched from the origin server. As noted in [1] and section 4.1.2, the concept of data-bins in JPIP are streaming-centric in contrast to the file-centric box structure of the basic JP2 file format. So, when all data-bins of a particular image have been received, we do not write out the image data as a regular codestream, whether in a raw format or encapsulated in a JP2 file format. This makes sense since the proxy will only ever interact with the image data in its data-bin form, so this facilitates loading the image data whenever it is requested by a client.

4.1.7 Signalling the completion of a server response

If all the available data in the cache for the requested image have been delivered to the client, and the connection with the origin server has been terminated, we use the window done end-of-response (EOR) terminator to signify that the window is partially complete. We allow the cache object to recover from a connection failure and attempt to reconnect indefinitely, so that the proxy does not miss any opportunity to fill its cache to service any pending client requests in the event the condition which caused the failure is remedied. If the proxy successfully reconnects, it resumes filling its cache with the new data. This particular EOR terminator is most appropriate amongst the terminators defined in [5], [9], but a side-effect of this is some clients may not be aware that the proxy has indeed updated its cache, unless the client generates a fresh request by changing its window of interest for instance.

4.2 Server policies

We describe below two different serving policies which may be adopted by the proxy server.

4.2.1 Packet forwarding

In the first policy, the proxy server acts as a packet forwarder, any and all packets which arrive at the proxy are simply cached and forwarded to the ultimate client. The proxy server ends up being completely transparent. This policy is straightforward if the proxy is empty, or it has served its entire contents relevant to the client window to the ultimate client. To generalize the serving policy to the case where there is existing data which have not already been sent to the ultimate client, we maintain a list of packets which have already been cached. The advantage of using this method is, in the case of packet forwarding, the only latency is due to network conditions and the delay involved in sending the packet from the origin server to the proxy, and from the proxy to the ultimate client. Keeping track of the order of packets received also allows the proxy to mimic the serving policy of the origin server. The disadvantage of course, is that a sub-optimal origin server policy results in a sub-optimal proxy server policy, in terms of the distortion of the image received by the ultimate client. This is perhaps acceptable if the origin server adheres to a formal progression order prescribed by the JPEG2000 standard, other than the LRCP progression. However, we cannot possibly predict the server policy in advance, although a server would certainly be better off following one of the prescribed progression orders in the first place. Thus packet forwarding is conservative in the sense that it potentially trades off optimality in the distortion sense for minimal delay.

4.2.2 Rate-distortion optimized policy

In the second policy, the proxy server keeps track of the list of packets cached in the proxy, as in the first policy. It then sequences these packets within a predetermined time window. In the EBCOT paradigm, packets are formed based on distortion values calculated during compression. The distortion values and truncation points for code-blocks contributing to a packet in a particular precinct are not available once the compressed codestream is generated. However, the distortion-length slope thresholds T_q which determine the truncation lengths at which an embedded code-block bit-stream is included in layer L_q, can be made available by including them in the COM marker segment of the codestream. In [3], a factor w_c is introduced which represents the factor by which the reconstructed samples from the packet overlap the window of interest. The quality layers L_q optimize the perceived image distortion for the entire image. In a window of interest, weighting precinct packets by multiplying T_q by w_c is a better indication of the reduction in distortion contributed by a packet in precinct c.

Within the predetermined window, the proxy accumulates packets from the origin server. Each time a packet arrives, the list of packets is resequenced according to their respective $w_c T_q$. By accumulating packets, obviously there is an enhanced delay due to the time window imposed. The larger the window, the more time we have to accumulate packets and thus the bigger the possibility that the distortion is reduced maximally for the number of packets transmitted to the ultimate client. However, the delay will also be larger. If the window is small enough, the policy defaults to the previous policy of packet forwarding. By artificially imposing a window on incoming or preexisting packets, we are indirectly imposing a length constraint on the distortion reduction available. Then we wish to minimize the distortion caused by the packets that have been accumulated.

As the packets always contribute a positive reduction in distortion, we just need to send as many packets as has been accumulated during the window period. In [3], the concept of resequenced packets is introduced, where packets are arranged in descending order of $w_c T_q$. Rather than resequence packets though, resequenced layers are used instead, with each layer containing at most one packet from each precinct. This is intended to minimize the number of empty packets created when the codestream is reconstructed dynamically. However JPIP does away with the need for this with the use of independent messages, where the message body is itself a valid self-contained JPEG2000 codestream. So packets can be sent separately with the overhead contribution from the JPIP message instead of the much larger overhead from empty packets in a resequenced layer. We thus reutilize resequenced packets in our proxy server. By sorting the packets and ensuring that

$$(w_c T_q)_1 \geqslant (w_c T_q)_2 \geqslant \ldots \geqslant (w_c T_q)_n \tag{4.1}$$

where n is the number of accumulated packets considered during our time window, minimizing the distortion out of a collection of packets can thus be achieved. This is equivalent to selecting the sequence s_{opt} out of all possible sequences S of JPIP packets which minimizes the distortion,

$$s_{opt} = \arg\min_{s \in S} D \tag{4.2}$$

where D is the distortion from the accumulated packets. Equivalently, we wish to minimize the function

$$\Lambda = D + \lambda L \tag{4.3}$$

subject to

$$L - n = 0 \tag{4.4}$$

where L is the number of accumulated packets.

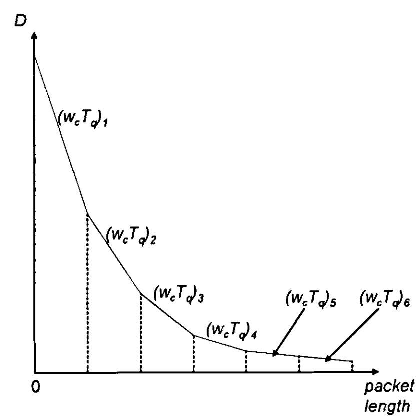

Figure 4.3: Optimal rate-distortion curve, packets sorted according to weighted thresholds

Graphically, the value of length L which minimizes D subject to the constraint on L is trivially the intersection of the line $L = n$ with the distortion curve. That is, send all n packets accumulated during the time window.

We make some assumptions about the network conditions: there is no packet loss, and we have a proxy-client bandwidth estimate of B. Let T be be the time duration of our window, including the time taken for packets to accumulate in the cache. Is there any relation between T and B? What is the effect on the proxy latency?

If the length of the packets accumulated is $\sum_{i=1}^{L} (l_{p_c})_i$ bytes, the bandwidth is B_t at time t, the expected delay is

$$\delta = \frac{\sum_{i=0}^{L} (l_{p_c})_i}{B_t} + T \tag{4.5a}$$

This is the delay of sending $\sum_{i=1}^{L}(l_{p_c})_i$ bytes formed from concatenating all the packets considered during the window. Would it be feasible or reasonable to increase the window duration T if the bandwidth B is large? Or do we vary the window duration according to the amount of packets/bytes received?

Consider the delay δ, letting

$$\delta = \frac{\sum_{i=1}^{L}(l_{p_c})_i}{B_t} + T \leqslant \delta_{\max} \tag{4.6}$$

The ratio $\frac{\sum_{i=1}^{L}(l_{pc})_i}{B_t}$ looks instructive:

- Case 1: If $\frac{\sum_{i=1}^{L}(l_{pc})_i}{B_t} << 1$, the proxy is receiving far less than it can transmit.
- Case 2: If $\frac{\sum_{i=1}^{L}(l_{pc})_i}{B_t} \approx 1$, the proxy is fairly relaying the packets it is receiving.
- Case 3: If $\frac{\sum_{i=1}^{L}(l_{pc})_i}{B_t} >> 1$, the proxy bandwidth is far exceeded.

In case 1, the cache has not that many packets in its cache, and whatever packets that are received can be rapidly sent. To minimize the distortion perceived by the client, we can probably afford to stretch out the time window, in order to receive more packets. Case 2 is slightly more ambiguous. There are a fair number of packets received, to work with when sorting for distortion optimality. Case 3 is obvious: for every iteration of the sending loop, sort the received packets and send right away, since the proxy cannot afford to wait.

Chapter 5

Results

5.1 Testing of Proxy Cache Performance

With the proxy server in place in the Telemedcare system as shown in Figure 4.1, we wish to investigate the performance benefit of the proxy server when multiple clients successively interrogate the proxy server for an identical window of interest in a test image. The test image parameters are:

- image size: 2944×1966 (width $\times$ height)
- image origin: (0,0)
- image is untiled
- number of components: 3
- component subsampling factors: (1,1), (1,1), (1,1)
- number of quality layers: 22
- number of DWT levels: 5

- precinct sizes: (64,64) (all resolutions)
- block size: (64,64)

Test image and region of interest

In this test, the first client sends a request from scratch for a window of interest with a size of 256×256 and offset of (1325,1063), for a duration of 45 seconds. Once the duration is complete, four other clients then successively request the same region for the same duration. At the end of the duration, the peak signal-to-noise ratio of the window of interest in the image retrieved is calculated. The PSNR of an image is defined as:

$$PSNR = 20\log_{10}\left(\frac{255}{\sqrt{MSE}}\right) \tag{5.1}$$

where 255 is the maximum pixel value for an image represented by 8 bits for every pixel in every colour component, and MSE is the mean squared error of the retrieved image with respect to the original image:

Client no.	PSNR	Bytes received (kB)
1	21.01	10.7
2	23.98	19.4
3	26.45	33.3
4	27.24	43.0
5	29.02	55.5

Table 5.1: Results for multiple clients successive browsing - 45s

Client no.	PSNR	Bytes received (kB)
1	24.50	20.6
2	27.06	39.2
3	29.11	55.4
4	35.16	89.8
5	∞	126.4

Table 5.2: Results for multiple clients successive browsing - 90s

$$MSE = \frac{1}{3mn} \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} \sum_{k=1}^{3} (A(i,j,k) - B(i,j,k))^2 \tag{5.2}$$

where $m \times n$ is the size of the image, A is the original image and B is the retrieved image, and 3 is the number of components (RGB).

The results are displayed in Table 5.1 and Figure 5.1:

The results clearly show that for the same browsing duration, each successive client retrieves an image region with a better PSNR than that for the previous client. If we perform the same test for a longer duration of 90 seconds, the results are shown in Table 5.2 and Figure 5.2:

The PSNR for the last client is essentially infinity, that is the region of interest has been brought to full quality for that client.

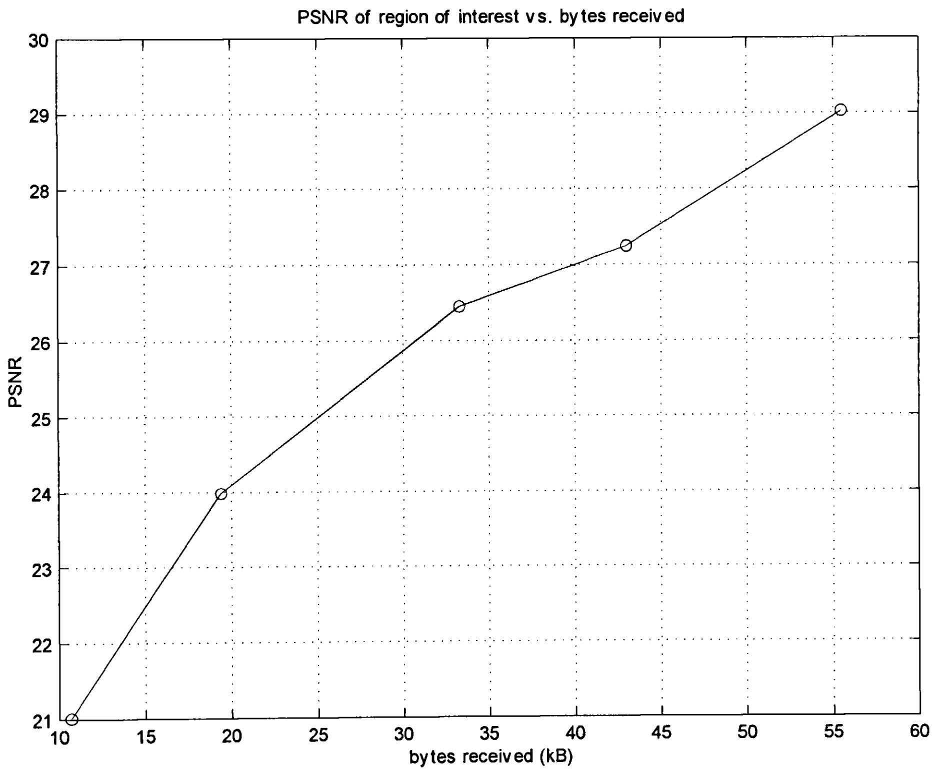

Figure 5.1: PSNR of region of interest vs. bytes received (duration of 45s)

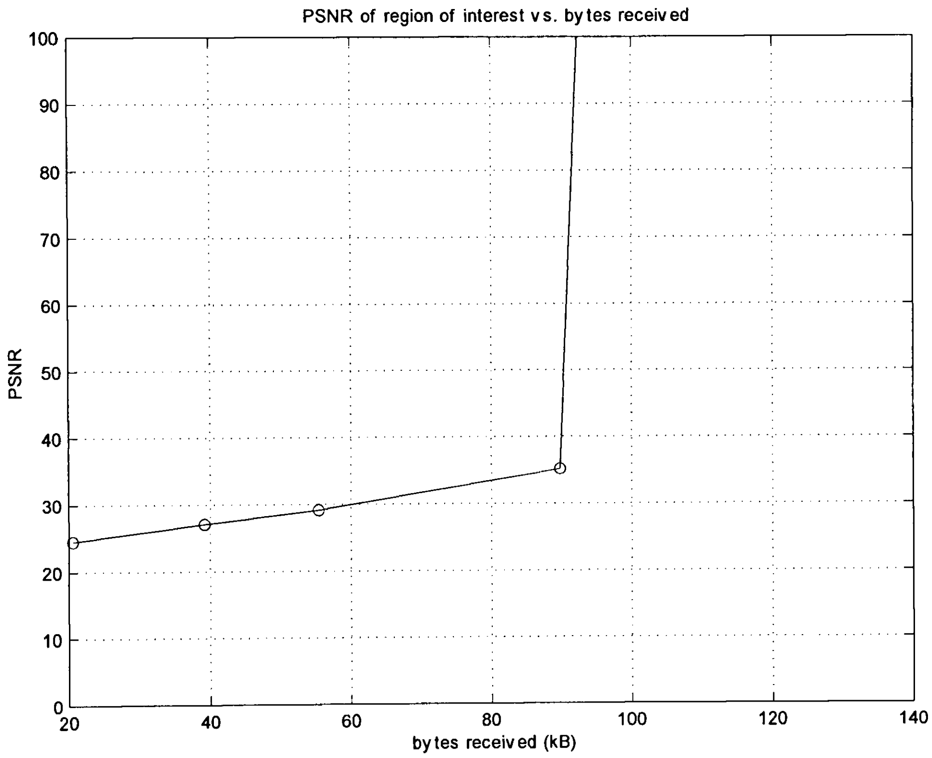

Figure 5.2: PSNR of region of interest vs. bytes received (duration of 90s)

5.2 Design of a Clinical Trial in a Nursing Home

In order to examine the feasibility of the image capture system, a clinical trial was designed for this purpose, but ultimately shelved (see section 5.3). The main goals were to gain an understanding of the following: the impact of integrating the system with other existing healthcare modules; optimal conditions for image acquisition; feedback on system convenience, ease of use and patient and doctor satisfaction. The knowledge obtained was to be used to improve on the current implementation to facilitate improved health outcomes.

The residential care facilities selected for this trial were the Montefiore Nursing Homes at Maroubra and Hunters Hill, Sydney. Discussions had taken place with senior staff at Montefiore Homes, and they had provided indications that they were happy to participate in the trial. However there was no formal exchange of letters with Montefiore Homes as this was pending upon the approval of the UNSW Human Research Ethics Committee (HREC). Once HREC approval was granted, the staff of Montefiore Homes would have presented the proposal and HREC approval for consideration by their own institutional ethics committee.

The study had its focus on care for elderly patients in a nursing home environment, and indeed elderly residents of the Montefiore residential care facilities are typically aged over 70. All residents in the homes at Maroubra and Hunters Hill would have been eligible to participate if they presented with lesions of clinical interest. Ethnicity was not relevant to the study. Subjects in the residential care facilities would not have had to perform any additional tasks outside of their normal care regime.

The methodology of the study was to be as follows. This is a qualitative study that seeks to evaluate the suitability of a data collection, storage and review format for digital images of clinical interest. All data and images obtained during the study are de-identified from any resident in the two nursing homes, but will be retained for the purpose of analysis and publication. However images may be used for additional research such as the development of algorithms to delineate wounds, identify gradations in texture and colour etc., for improved tracking of longitudinal changes. There would appear to be no particular need to destroy what would be a useful database of de-identified images that may prove useful for future research. If the HREC requires it however, the images would be destroyed at the completion of this specific project. Data and images obtained would be stored in secure servers located at the Biomedical Systems Laboratory at UNSW.

The recording of digital images and their review by the patients' own GP or another authorized member of the clinical care team represents a potential improvement on existing care by permitting a clinician to remotely view the images and advise on treatment. The outcome of research was expected to yield an improved visual diagnostic method (image capture via consumer digital cameras), which will be incorporated into existing telemedicine system developed by the researchers for use in residential care facilities or rural and remote community health centres.

5.3 Decision to Abort Clinical Trial in Nursing Home

Success of the trial depends on the convenience of the operation of the imaging system as well as the quality and performance. There were always concerns that tethering the camera to the PC would unnecessarily inconvenience the user (typically the nurse) and tend to make the imaging system less attractive. Very recently availability of wireless cameras supported by PTP effectively made the harnessed version of the system no longer competitive and therefore likely to be quickly replaced by the wireless LAN version. This will take some months and is beyond the scope of this thesis. Accordingly it was considered that given the cost and the effort required in mounting a clinical trial of the imaging system, it would be best to carry out only one clinical trial once the wireless LAN camera interface was designed, implemented and tested. Hence it was decided to abort the trial.

5.4 Evaluation of Clinician Response to Client-side Use of System

A repository of five sample images (see Appendix A) of approximately 3 megabytes each, and related to dermatology, was prepared for a panel of specialists to evaluate. The panel consists of a registered nurse, a medical doctor and an emergency specialist. A questionnaire was prepared beforehand. The questionnaire contains a list of questions which asks the respondents to either state their degree of agreement with a particular question, or alternatively to rate a statement from very poor to very good. At the end of the questionnaire there are a few questions which prompt the respondents to input some subjective opinions, possibly providing the opportunity to expand on questions previously asked. The questions are broadly divided into several categories: ease of use and performance, clinical use and benefits, and comments and suggestions.

In addition to the questionnaire, a fairly quantitative measurement was performed. Here, the set of five images were evaluated. The images were to be browsed on firstly a modem connection, and then a broadband connection. The following scenarios for a particular image were considered:

- Scenario A: The image was browsed as-is (i.e. without selecting a region of interest (ROI) with the focus box) on a modem connection. The time in seconds taken for the image to be received was recorded.
- Scenario B: The image was browsed on a modem connection, but a ROI was

Test Image No.	Scenario A	Scenario B	Scenario C	Scenario D
1	79	29	4.3	2.8
2	77	55	4.8	3.2
3	80	24	4.7	2.7
4	77	52	4.8	3.2
5	93	43	6.8	3.0

Table 5.3: Evaluation panel results - average times

selected. The time in seconds taken until the earliest instant a useful diagnosis can probably be made, was recorded.

- Scenario C: The image was browsed as-is (i.e. without selecting a ROI with the focus box) on a broadband connection. The time in seconds taken for the image to be received was recorded.
- Scenario D: The image was browsed on a broadband connection, but a ROI was selected. The time in seconds taken until the earliest instant a useful diagnosis can probably be made, was recorded.

The results of the evaluation questionnaires are extremely encouraging (Appendix C). In general, the respondents find the system and the various interfaces easy to use and operate. From the retrieved images, the respondents are happy with the image quality and are able to identify the presence of skin disorders as well as discern variations in skin color, shape, surface whenever the lesions were identified. All respondents felt that the system would be beneficial for use in healthcare in the community and in remote areas, and would provide quicker access to specialist knowledge thereby reducing treatment and waiting times for patients. The respondents agreed that they would be happy to use the system on a regular basis, albeit with a few suggestions on how to improve on the user interfaces.

The results for the quantitative tests are shown in Table 5.3, while plots of the average times are shown in Figs 5.3 and 5.4. The modem download speeds ranged from 4-6kB/s, while the broadband speeds were 256kbps (31kB/s) on ADSL connections.

Test Image No.	Scenario A	Scenario B	Scenario C	Scenario D
1	79	19	2.3	1.0
2	77	66	2.4	1.6
3	80	7.1	2.1	0.58
4	83	63	2.4	1.6
5	65	26	3.5	1.0

Table 5.4: Evaluation panel results - standard deviation of times

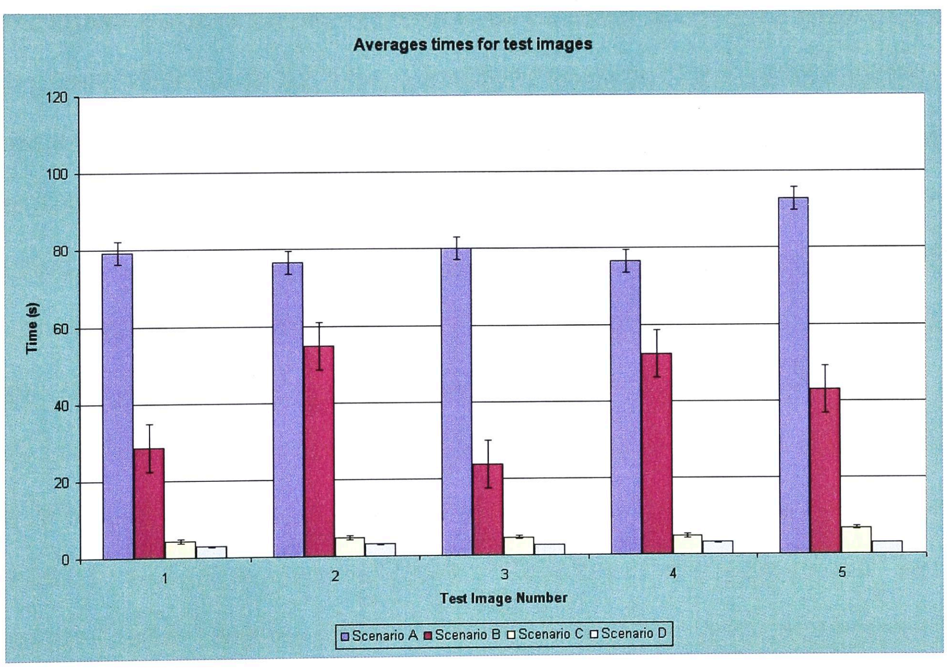

Figure 5.3: Average times for test images

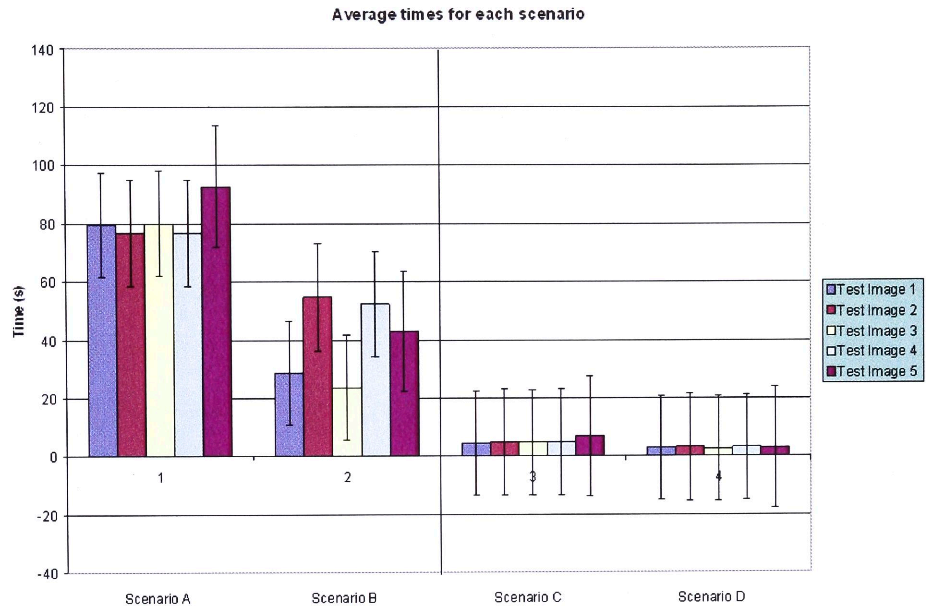

Figure 5.4: Average times for each scenario

From Figures 5.3 and 5.4, as expected, the retrieval times for a modem connection are much greater than those for a broadband connection. For a modem connection, the earliest times before a diagnosis is possible, with a ROI selected, range from about a third to three-quarters of the times for the images to be browsed as-is without a ROI selected. For a broadband connection, the time differences are almost negligible. Clearly the benefits in terms of reduced browsing times are more apparent for a modem connection; this is alluded to by one respondent who stated that the ROI selection capability is a time-saver for modem connections but not so much for broadband connections. When using a broadband connection, we can probably afford to wait for the image without selecting a region of interest before making a diagnosis. Test images 2 and 4 require more time for retrieval before diagnosis is possible on a modem connection. Referring to Appendix A, test images 2 and 4 have relatively large regions of interests, occupying much of the central region of the image. Test image 5 also has a large region of interest, but respondents tend to select the smaller left hand side of the region of interest, yielding a smaller time before diagnosis is possible. The error bars for the times before diagnosis for a modem connection are rather large, but this can be explained by the variability of the connection speeds, and the subjectivity of the respondent with regard to the region which is of interest and its quality for diagnosis.

The results indicate that it is possible to scale the size of the images stored at the server upward, without losing too much performance in terms of download time, even on a modem connection. For instance, referring to the results for test image 3 in Table 5.3, supposing the source image is increased in size by a factor of 10, a rough estimate of the times taken to browse the image as-is and for the earliest diagnosis with ROI on a modem connection would be 1025s and 205s respectively. However, the time for earliest diagnosis does not necessarily have to scale linearly. By adding more quality layers and assigning them bitrates properly, this time may possibly be reduced. Increasing the size of the images allows more detail and resolution to be added, which may further ease concerns about the quality of images deemed acceptable for medical diagnosis purposes.

For a store-and-forward tele-imaging system, the time taken before a GP may make a diagnosis will obviously be the time taken for the entire image to be downloaded. Thus we may treat the time taken for the image to be received as mentioned above to be an informal guide as to how a store-and-forward system would have performed. Allowing a GP to browse images progressively by quality seems to provide some saving in time when compared to waiting for images to completely download before making a diagnosis.

Chapter 6

Conclusion and Future Work

In this project, we endeavoured to develop an imaging system which incorporates consumer camera technologies and JPEG2000 technology, in conjunction with the Telemedcare telehealth system in use in the Biomedical Systems Laboratory. A workable system has been put into place and is currently integrated into the Telemedcare system in a test capacity. The system is placed at a location which has a Telemedcare medications trolley and with the appropriate infrastructure in place, a small evaluation took place with the other medication management and clinical measurements systems.

The home user is able to capture images as JPEG2000 using a commercial-off-the-shelf imaging device and imaging software, transfer them via a connection to the Telemedcare Home Desktop, and then leverage the existing system to transmit these images to the central database server. Once uploaded, the images can be browsed remotely by a client through an interactive application, with the images improving with quality as more data arrives from the server. As part of the imaging system, an experimental JPIP proxy server was developed, to improve the efficiency of the image browsing experience. Issues regarding the development of such a proxy were explored. Among these are a description of the structure of the proxy server, image cache data management, the serving of client requests by the proxy, the serving of meta-data, the concepts of source and source clients, assignment of target-IDs, and the signalling of a completed proxy server response. In addition, a few serving policies were discussed although they were not completely implemented in this project.

It has been found that clients are able to retrieve images captured by a particular user simply and effectively. Based on the evaluation panel, they were happy with image quality, and were able to discern and distinguish between skin disorders of various kinds. The system was judged to have the potential of granting faster access to specialist knowledge, which would benefit patients at health facilities by reducing treatment and waiting times.

Many areas of research and development can be envisioned to broaden the Telemedcare imaging system and experience.

- Client application improvements can be added. These naturally include standard image processing functions (adjusting brightness and contrast, blurring and sharpening, filtering, morphology etc.), and adding a facility for annotations and labelling of images.
- A smart annotation system which possibly utilizing XML, network and database technologies and the JPEG2000 file format XML container. Annotations could be tied in with image metadata and hence the proper and efficient ways of serving these metadata.
- Image retrieval and classification systems, to facilitate sorting and searching of images in storage.
- Improved proxy support: improved metadata, JPX file format support, HTTP-TCP transport, cache management, dynamic bandwidth allocation when there are multiple clients.
- Optimal camera imaging conditions and parameters in real nursing or home care situations.

- Extension of AcquirePhoto module's viewfinder to support real-time transmission of images, like a live video stream perhaps using Motion JPEG2000.
- Investigation into optimal coding parameters for use when transcoding downloaded camera image to JPEG2000. E.g. code-block size, number of layers, number of resolutions, bit-rate allocations, ROI designations, etc..
- Extension of the image acquisition program to incorporate PTP-IP. PTP-IP is an extension of PTP to facilitate the use of PTP over wireless LAN networks. This would allow the use of 1st generation wireless LAN cameras released recently.

Bibliography

[1] D. Taubman and R. Prandolini, *Architecture, Philosophy and Performance of JPIP: Internet Protocol Standard for JPEG2000.* Visual Communications and Image Processing 2003. Proceedings of the SPIE, Volume 5150, pp. 791-805, 2003

[2] D. Taubman and M. W. Marcellin, *JPEG2000. Image Compression Fundamentals, Standards and Practice*, Kluwer Academic Publishers, 2002.

[3] D. Taubman and R. Rosenbaum, *Rate-distortion Optimized Interactive Browsing of JPEG2000 Images*, IEEE International Conference on Image Processing (ICIP2003), volume 3, pp. 765-768, September 2003.

[4] J. P. Ortiz, V. G. Ruiz and I. García, *Improved JPIP Protocol with Proxy Caching.* In I. García, L.G. Casado, V.G. Ruiz, and J.J. Fernández, editors, Actas de las XV Jornadas de Paralelismo. Computación de Altas Prestaciones, pages 322-327, Almería, España. Servicio de publicaciones de la Universidad de Almería. Via http://www.ace.ual.es/~vruiz/investigacion/node7.html. September 2004.

[5] D. Taubman, *Proposal and Implementation of JPIP (JPEG2000 Internet Protocol) in Kakadu V3.3.* Via http://www.kakadusoftware.com. August 2002.

[6] Kakadu Software by David Taubman. http://www.kakadusoftware.com

[7] ISO/IEC 15444, *Information Technology - JPEG2000 Image Coding System - Part 1: Core coding system*, 2004.

[8] ISO/IEC 15444, *Information Technology - JPEG2000 Image Coding System - Part 2: Extensions,* 2004.

[9] ISO/IEC 15444, *Information Technology - JPEG2000 Image Coding System - Part 9: Interactive tools, APIs and protocols,* 2006.

[10] N. H. Lovell, F. Magrabi, B. G. Celler, K. Huynh and H. Garsden, *Web-based Acquisition, Storage, and Retrieval of Biomedical Signals.* IEEE Engineering in Medicine and Biology, May/June 2001, pp. 38-44.

[11] B. G. Celler, N. H. Lovell, J. Basilakis, F. Magrabi and M. Mathie, *Home Telecare for Chronic Disease Management.* Proceedings of the 23rd Annual EMBS International Conference, October 25-28 2001.

[12] N. H. Lovell, B. G. Celler, J. Basilakis, F. Magrabi, K. Huynh and M. Mathie, *Managing Chronic Disease with Home Telecare: A System Architecture and Case Study.* Proceedings of the Second Joint EMBS/BMES Conference, Houston TX, October 23-26 2001.

[13] H. K. Huang, *PACS and Imaging Informatics: Basic Principles and Applications.* John Wiley and Sons, 2004.

[14] Australian Institute of Health and Welfare, *Residential Aged Care in Australia 2003-04, A Statistical Overview.* Canberra, 2005.

[15] Aware, Inc. http://www.aware.com

[16] Integrated Distributed Environment for Application Services in e-Health (IDEAS in e-Health). http://www.ideas-ehealth.upv.es

[17] Edward B. Suh, S. Warach, H. Cheung, S. A. Wang, P. Tangiral, M. Luby, R. L. Martino, *Web-based medical image archive system.* Proc. SPIE Vol. 4685, Medical Imaging 2002, pp. 31-41.

[18] E. S. Lee, I. S. Kim, J. S. Choi, B. W. Yeom, , H. K. KIM, G. H. Ahn, S. Y. Leong, *Practical Telepathology Using a Digital Camera and the Internet.* Telemedicine Journal and e-Health, vol. 8, no. 2, 2002.

[19] I. Maglogiannis, *Design and Implementation of a Calibrated Store and Forward Imaging System for Teledermatology.* Journal of Medical Systems, vol. 28, no. 5, October 2004.

[20] LibUsb-Win32. http://libusb-win32.sourceforge.net

[21] Libptp2 Project Page. http://libptp.sourceforge.net

[22] PIMA 15740:2000, *Photography – Electronic still picture imaging – Picture Transfer Protocol (PTP) for Digital Still Photography Devices.* Via http://www.i3a.org

[23] CIPA DC-005 (PTP-IP). Via http://www.cipa.jp/ptp-ip/index_e.html.

[24] ColdFusion MX7. http://www.adobe.com/products/coldfusion

[25] J. D. Whited, *Teledermatology. Current status and future directions.* Am J Clin Dermatol. 2001;2(2):59-64.

[26] M. S. Houston, J. D. Myers, S. P. Levens, M. T. McEvoy, S. A. Smith, B. K. Khandheria, W. K. Shen, M. E. Torchia and D. J. Berry, *Clinical consultations using store-and-forward telemedicine technology.* Mayo Clin Proc. 1999 Aug;74(8):764-9.

[27] E. A. Krupinski, B. LeSueur, L. Ellsworth, N. Levine, R. Hansen, N. Silvis, P. Sarantopoulos, P. Hite, J. Wurzel, R. S. Weinstein and A. M. Lopez, *Diagnostic*

accuracy and image quality using a digital camera for teledermatology. Telemed J. 1999 Fall;5(3):257-63.

[28] R. X. Murphy, Jr., M. A. Bain, T. E. Wasser, E. Wilson, W. J. Okunski, *The Reliability of Digital Imaging in the Remote Assessment of Wounds: Defining a Standard.* Annals of Plastic Surgery, vol. 56, no. 4, pp 431-436, April 2006.

[29] D. J. Eedy and R. Wootton, *Teledermatology: a review.* British Journal of Dermatology 2001 144: pp. 696-707.

[30] Independent JPEG Group. http://www.ijg.org

[31] P. Corr, I. Couper and S. J. Beningfield, *A simple telemedicine system using a digital camera.* Journal of Telemedicine and Telecare, Aug 2000, vol. 6, no. 4, pp. 233-236(4). *Publisher: Royal Society of Medicine Press*

[32] A. Marco, M. G. Albanesi, *Extracting Metadata fron JPEG2000 Compressed Images Using Web Services.* Proceedings of the Seventh International Workshop on Computer Architecture for Machine Perception (CAMP'05) 2005, pp. 315-320.

[33] E. R. Harold, W. S. Means, *XML In A Nutshell, 3rd Edition.* O'Reilly, 2004.

[34] N. Weiss, *Utilizing XML and XSLT in ColdFusion MX.* Via http://www.macromedia.com/devnet/mx/coldfusion/articles/xmlxslt.pdf.

Appendix A

Test Images

Figure A.1: Test image 1

Figure A.2: Test image 2

Figure A.3: Test image 3

Figure A.4: Test image 4

Figure A.5: Test image 5

Appendix B

Submitted Conference Papers

A Novel Image Capture System for Use in Telehealth Applications

Ed W. Lim, Branko G. Celler, Jim Basilakis and David Taubman

Abstract— **A novel image capture and retrieval system has been developed for use in a range of telehealth applications in the home and in residential care facilities. The system is based around the JPEG 2000 standard and uses the PTP protocol for image capture from any high resolution digital camera and the Kakadu suite of JPEG2000 utilities to serve the collected images via a proxy server over any available communication channel from telephone lines to broadband services. When coupled with an image processing system such as the AMWIS system for pressure wound management, the system provides a high level of clinical functionality suitable for a wide range of telemedicine applications in rural and remote sites.**

I. INTRODUCTION

THIS paper is concerned with the development of a low-cost imaging system for use in the UNSW home telecare system or Telemedcare. The imaging system allows a patient at a private home or nursing care facility to capture an image with a consumer digital camera, and deliver the image efficiently over the Internet to the medical practitioner in charge of the patient. The emphasis is on the acquisition and delivery of images from patient to practitioner. The system is thus potentially a good complement to a dedicated wound imaging system like the AMWIS system being employed in several sites around Australia [18]. The AMWIS system allows a health officer to accurately measure the area of a wound and produce a comprehensive wound record to be presented to a specialist. However, acquiring and inserting an image into the system has to be performed manually, and transmission of an image is done via email, which may be inefficient for large images.

II. TELEMEDCARE

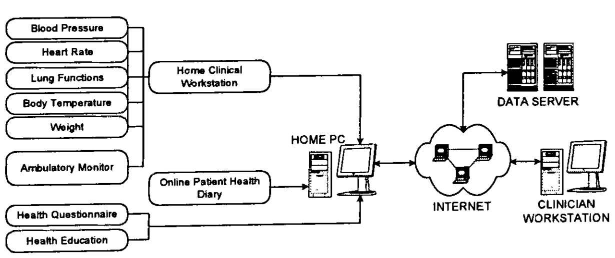

Figure 1-Diagram of Telemedcare System

The in-house home telecare system, or Telemedcare [7], [8], [9], provides a home patient with self-administered functional health status measures in the form of questionnaires, remote monitoring of parameters of daily changes to patient status, and clinical measuring devices. Acquisition of biomedical data or questionnaire responses is carried out via the Internet. Data is stored in a database and later analyzed for longitudinal trends. This data is available on request to the client (usually a general practitioner or GP) via a web browser.

The patient is provided with a clinical workstation. It is modular and includes:

1) ECG, blood pressure, spirometer, blood oximetry, sphygmomanometer.
2) Low power RF interface to ambulatory patient worn device for telephone voice connection on emergency button press & triaxial accelerometer.
3) Accurate patient weight and temperature.
4) Ambient temperature, light and humidity.

Based on case studies and clinical trials, the Telemedcare system has been shown to be functional and effective. It has been deployed in actual home settings and is a useful tool for the management of chronic illnesses like congestive heart failure and chronic obstructive pulmonary disease [8]. For home telecare, the minimization of cost is important, and the home telecare system achieves this by emphasizing modularity in its design, so that the modules required can be tailored to suit individual patients' needs. The patient user interfaces are designed to be simple and intuitive, so that users of the system need not have had any prior experience using computers.

III. IMAGING MODALITIES FOR TELECARE SYSTEMS

A. *Common modalities*

The imaging modalities that are most commonly used in telemedicine are the `store-and-forward system' modality [14], [15], [17], where images are saved locally, then forwarded by email to the recipient, and the real-time video/web conferencing modality [10], [14], [17]. For both modalities, image transmission over a network with limited bandwidth hampers their effectiveness. In the former, the GP may need to wait for a long time for large images to arrive. In the latter, the conferencing session may experience high delay and latency. In addition, conferencing systems tend to incur higher implementation costs.

B. *JPEG2000 and Progressive Transmission*

Telemedcare does not currently have any facility for imaging acquisition and analysis. This project attempts to utilize progressive image transmission as an imaging modality for Telemedcare. We use the JPEG2000 imaging standard which allows for efficient compression, scalability

E.W. Lim, B.G. Celler, Jim Basilakis and D. Taubman are with the University of New South Wales, Sydney 2052 Australia (e-mail: contact ed.lim@student.unsw.edu.au,, or b.celler@unsw.edu.au)

and codestream flexibility amongst its many benefits [2], [5]. The standard allows progressive transmission of compressed image based on precinct, component, resolution or layer. The standard also defines the JPEG2000 interactive protocol (JPIP), which is based on HTTP, for efficient serving of images over data networks [1], [3], [6].

With JPIP we can efficiently browse (or serve) large images remotely, and with a judicious choice of progression we can browse images with good quality even over low datarate networks. For instance by encoding an image with a number of quality layers and browsing this image remotely, a low quality or grainy representation of the image can firstly be brought to view. As additional data is sent the image quality is successively improved. In addition, a region of interest of the image may be selected such that only that region is improved in quality. Quality progression allows the user to have a `view' of the image even at low quality instead of waiting for the image to load from top to bottom as is conventionally done, and region-of-interest selection allows the user to enhance a particular location of the image while ignoring its surroundings which may be deemed to be unimportant.

C. *Use of COTs imaging devices in telemedicine*

The use of digital cameras in telemedicine is not new. In [16], such cameras used in the context of teledermatology were found to provide clinicians with high-quality images suitable for diagnosis, with outcomes comparing well with in-person consultations.

IV. System Overview

Digital cameras are commonly available with a proliferation of different models from many vendors, but we need to interface with the camera firmware to develop custom software for the system. To achieve this, we use an open standard called Picture Transfer Protocol (PTP) or formally known as PIMA/ISO 15740 [13]. PTP defines a set of containers and data formats, standard imaging referencing behaviour, operations, responses, events, device properties and datasets for transferring data to and from a digital still photography device and devices such as computers and printers. PTP also defines optional operations and formats, and mechanisms for extending the protocol by individual vendors. This imaging protocol removes the need for camera-specific drivers. PTP supports common device controls like image capture and image download, providing us with flexibility in application development. The PTP standard is also transport-independent, with details needed to use PTP with USB provided with the standard. The use of PTP is primarily advantageous because it eliminates the need to use proprietary protocols for different device vendors to communicate with the camera firmware. All major camera vendors have provided support for PTP in most (if not all) of their products since 2002 [12]. However we are still subject to the limitation of the subset of PTP operations and properties that manufacturers choose to support in their camera firmware.

The most prevalent image format in use with imaging devices is JPEG, and thus we will be working primarily with this format. In order to exploit the benefits of JPEG2000, the image capture program incorporates means for converting image file formats supported by digital cameras to JPEG2000. Commonly used web browsers like Netscape, Firefox and Internet Explorer have in-built capabilities to render and display JPEG images. During the transcoding operation, we also use the JPEG converter to produce a low-resolution thumbnail of the original image captured by the imaging device. We can thus take advantage of the browser capabilities to easily present previews of the images stored in the central database to the client.

The Telemedcare system uses a hybrid system for storage of biomedical signals acquired from the clinical workstation. In a hybrid system, a shared-file database collects data locally, then periodically synchronizes this data with the central database server (an umbrella term for processing, broker and database servers) through the Internet or a local area network. In this way the system is able to utilize network resources more efficiently and support multiple users [7]. The central database itself stores the clinical data, including ECG and blood pressure readings.

The use of JPEG thumbnails allows the preview of many images at once with only a small cost in bandwidth. There is thus no need to send the full JPEG2000 image if it is not required. After synchronization the processing server will open a connection to the database server, process the image records and store the main JPEG2000 images in a separate directory on the processing server, while the thumbnails will be reinserted into a separate table as part of the database server. Other metadata related to the images can be inserted together with the thumbnails as part of a comprehensive collection of information related to the patient images generated.

Clients may be served by a proxy server instead of the central database, which behaves as the ultimate server for all the images collected from the patients. The client application establishes a stateful JPIP session with the server. The proxy is JPIP-compliant, i.e. it accepts JPIP requests from a client and issues its own JPIP requests to the ultimate server. The proxy caches images, which it requests, which it can then use to satisfy subsequent requests for the same images, possibly from different clients. The proxy can cache and serve incomplete images. We adhere to a JPIP principle that the server may serve image data in any order it pleases, thus the proxy just serves whatever is present in the image cache. The proxy server however is deficient in the following senses:

1) the proxy is filled with image data only when prompted by client request.
2) it only supports HTTP transport although JPIP provides for HTTP-TCP as well
3) it has only been tested with meta information necessary

for opening and displaying the JP2 file format.

We use open source software for PTP and USB available via [11] and [12]. To work with JPEG2000, we use the Kakadu Software [4]. Kakadu is a comprehensive, optimized and fully compliant software for developing JPEG2000 applications. Kakadu is not freely available, but a license can be obtained at an affordable cost, and the license grants the user with unlimited upgrades and compilations.

Each of the proxy and client applications uses or is based on Kakadu. For the patient-side camera capture utility, we use the open source LibUSB and LibPTP libraries and a LibUSB wrapper to allow custom USB communication. If the platform is non-Windows XP an additional open-source generic digital camera driver must also be installed in order to talk to the imaging device. We use Cold Fusion markup and SQL for database queries as noted above. The applications themselves have been executed on Windows 2000 and XP platforms. A precompiled Kakadu DLL is required to run the programs.

There are no special hardware requirements for running the applications, except for a personal computer with dial-up or broadband internet access, as well as a digital camera for image capture which supports PTP.

V. System Operation

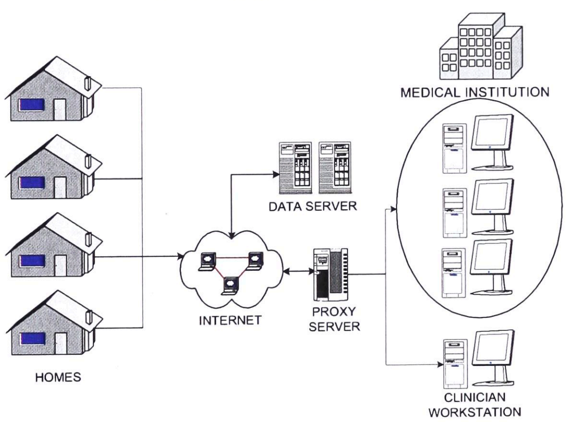

Figure 2 - Structure of Imaging System

A. *Home patient-side description*

At the patient side, the home clinical workstation presents a simple user interface to the patient. The patient captures an image using a digital camera connected to the workstation. At present the transport medium used is USB. The image is captured as JPEG but converted into JPEG2000. A JPEG thumbnail is created simultaneously. The image is encoded by Multipurpose Internet Mail Extensions (MIME) encoding and sent as XML to the Telemedcare database, along with the thumbnail. The use of MIME and XML is to conform with the existing data encapsulation paradigm in use in Telemedcare. The thumbnail is stored in the database along with other clinical data produced by the home workstation. The actual image is stored separate from the main database, without modifications.

The image acquisition software is set up by default for macro or close-up shots. Advanced users may access a menu and adjust parameters to suit their preferences, including flash, white balance, exposure, zoom, aperture and shutter speed settings, image quality and file sizes. Available parameters are determined at run-time by querying the camera, as not all parameters may be supported by every camera

The existing Home Desktop software infrastructure is used to perform synchronization with the central database server and transmit the image data over the network.

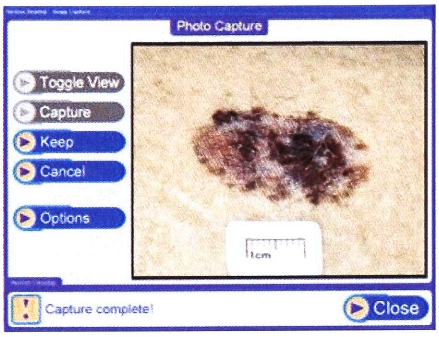

Figure 3 - Image Capture User Interface

B. *Client-side description*

At the client or GP workstation, the client browses images available on the central database through a web interface. The database returns the thumbnail previews for the current patient. When the client clicks on the preview, an application is launched which allows interactive browsing of the requested image via JPIP. The client may be served by a proxy cache, but this would be transparent to the client.

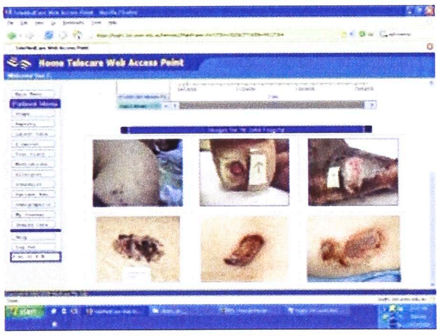

Figure 4 - Client-side Interface

C. *Database description*

During episodic synchronization with the patient workstations, the central database server accepts incoming “suitcase” files containing the biomedical patient data and images, together with the thumbnail previews, and inserts them into a table. All newly-accepted data files will be stored as records in this table, including our image data. The record which appears in the table remains the compressed data encased in XML format. So, we need to reverse all the processes which were involved in preparing the image data for transmission on the Home Desktop side. The processing server runs a Cold Fusion server which parses the web requests from clients, and accesses the database server accordingly. The actual image requested is served by a JPIP server running separately.

D. *Interactive application description*

The interactive application is based on the *kdu_show* application provided with Kakadu to view JPEG2000 images. As well as able to display local or remote images within a selected region-of-interest, *kdu_show* is also able to edit optional metadata such as labels in the images, and perform simple spatial manipulations like rotate and flip.

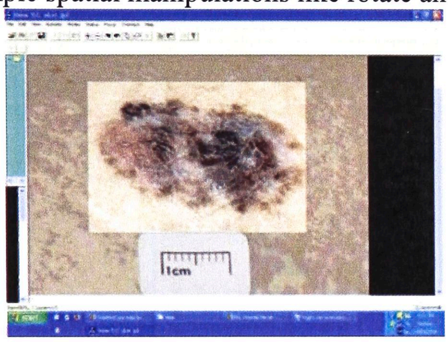

Figure 5 - Interactive application

VI. Mini-trial and Results

Sample images as shown in Figures 3-5 were taken at a number of locations and transmitted via the system to demonstrate the system's usability. A Canon Powershot Pro1 was used to acquire the images. The following settings were used for the camera: F8.0 aperture (smallest), 2"5 shutter speed, superfine image quality, large image size, macro mode, auto white balance, no flash. Aperture and shutter speed were found to be particularly important, where a combination of small aperture and slow shutter speed yielded sharpest images and maximum depth-of-field. The slow shutter speed however necessitated a steady camera while an image was shot.

Preliminary work has also been completed to link the imaging system to the Alfred Medseed Wound Imaging System V1.0 [18] which will permit a baseline health assessment, and ongoing AMWIS wound assessment for duration of wound management to assess the base-line healing rate (% change in wound size per week).

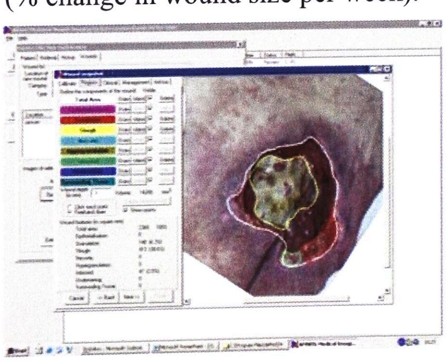

Figure 6. AMWIS analysis of wound profile

Future Work

Currently the imaging system leverages the existing infrastructure to transmit images between the Home Desktop and the central database server. This was a design decision such that the imaging system could fit into Telemedcare easily. However this link is not optimized for large data like images, and will be subject to further research. Also, we intend to establish an image indexing and retrieval system, possibly separate from the current database implementation. We plan to add two major forms of functionality to the application: simple image processing tools and the ability to annotate the images displayed using the application. Wireless LAN cameras have recently been released in the market; support for these may be added in future.

References

[1] D. Taubman and R. Prandolini, "Architecture, Philosophy and Performance of JPIP: Internet Protocol Standard for JPEG2000," Visual Communications and Image Processing 2003, in *Proceedings of the SPIE*, vol. 5150, pp. 791-805, 2003.

[2] D. Taubman and M. W. Marcellin, *JPEG2000. Image Compression Fundamentals, Standards and Practice.* Kluwer Academic Publishers, 2002.

[3] D. Taubman, "Proposal and Implementation of JPIP (JPEG2000 Internet Protocol) in Kakadu V3.3". Available: http://www.kakadusoftware.com. August 2002.

[4] Kakadu Software by David Taubman. Available: http://www.kakadusoftware.com.

[5] ISO/IEC 15444, *Information Technology - JPEG2000 Image Coding System - Part 1: Core coding system.* 2000.

[6] ISO/IEC 15444, *Information Technology - JPEG2000 Image Coding System - Part 9: Interactive tools, APIs and protocols.* 2003.

[7] N. H. Lovell, F. Magrabi, B. G. Celler, K. Huynh and H. Garsden, "Web-based Acquisition, Storage, and Retrieval of Biomedical Signals," *IEEE Engineering in Medicine and Biology*, pp. 38-44, May/June 2001.

[8] B. G. Celler, N. H. Lovell, J. Basilakis, F. Magrabi, M. Mathie, "Home Telecare for Chronic Disease Management," in *Proceedings of the 23rd Annual EMBS International Conference*, October 25-28 2001.

[9] N. H. Lovell, B. G. Celler, J. Basilakis, F. Magrabi, K. Huynh, M. Mathie, "Managing Chronic Disease with Home Telecare: A System Architecture and Case Study," in *Proceedings of the Second Joint EMBS/BMES Conference*, Houston TX, October 23-26 2001.

[10] H. K. Huang, *PACS and Imaging Informatics: Basic Principles and Applications.* John Wiley and Sons, 2004.

[11] LibUsb-Win32. Available: http://libusb-win32.sourceforge.net.

[12] Libptp2 Project Page. Available: http://libptp.sourceforge.net.

[13] PIMA 15740:2000, *Photography -- Electronic still picture imaging -- Picture Transfer Protocol (PTP) for Digital Still Photography Devices.* Available: http://www.i3a.org.

[14] J. D. Whited, "Teledermatology. Current status and future directions," *Am J Clin Dermatol*, vol. 2(2), pp. 59-64, 2001.

[15] M. S. Houston, J. D. Myers, S. P. Levens, M. T. McEvoy, S. A. Smith, B. K. Khandheria, W. K. Shen, M. E. Torchia, D. J. Berry, "Clinical consultations using store-and-forward telemedicine technology," *Mayo Clin Proc.*, vol. 74(8), pp. 764-769, August 1999

[16] E. A. Krupinski, B. LeSueur, L. Ellsworth, N. Levine, R. Hansen, N. Silvis, P. Sarantopoulos, P. Hite, J. Wurzel, R. S. Weinstein, A. M. Lopez, "Diagnostic accuracy and image quality using a digital camera for teledermatology," *Telemed J.*, vol. 5(3), pp. 257-263, 1999.

[17] D. J. Eedy, R. Wootton, "Teledermatology: a review," *British Journal of Dermatology*, vol. 144, pp. 696-707, 2001

[18] N. Santamaria, L. Clayton, "Cleaning Up. The Development of the Alfred/Medseed Wound Imaging System," *Collegian*, vol. 7(4), pp. 14-15, 17-18, October 2001

Appendix C

Evaluation Panel Questionnaire

EVALUATION QUESTIONNAIRE FOR IMAGING SYSTEM

Section 1 – Evaluation of Images

In this section, a set of five images are to be evaluated. The images are to be browsed on firstly a modem connection, and then a broadband connection. Firstly, the image is browsed as-is (i.e. without selecting a region of interest with the focus box) on a modem connection. The time taken for the image to be received is recorded. Next, the same image is browsed, but a region of interest is selected. The time taken until the *earliest instant that a useful diagnosis can probably be made*, is recorded. These steps are repeated for a broadband connection.

<u>IMAGE 1</u>
1) Modem Connection
Time for complete download:

....

Time for earliest diagnosis:

....

2) Broadband Connection
Time for complete download:

....

Time for earliest diagnosis:

....

<u>IMAGE 2</u>
1) Modem Connection
Time for complete download:

....

Time for earliest diagnosis:

....

2) Broadband Connection
Time for complete download:

....

Time for earliest diagnosis:

....

IMAGE 3

1) Modem Connection
Time for complete download:

....

Time for earliest diagnosis:

....

2) Broadband Connection
Time for complete download:

....

Time for earliest diagnosis:

....

IMAGE 4

1) Modem Connection
Time for complete download:

....

Time for earliest diagnosis:

....

2) Broadband Connection
Time for complete download:

....

Time for earliest diagnosis:

....

IMAGE 5

1) Modem Connection
Time for complete download:

....

Time for earliest diagnosis:

....

2) Broadband Connection
Time for complete download:

....

Time for earliest diagnosis:

....

For Section 2 and 3, please tick the most appropriate response.

Section 2 – Ease of Use, Performance

1. How would you rate your level of computer expertise?

- ☐ very good
- ☐ good
- ☐ adequate
- ☐ poor
- ☐ very poor

2. Based on your response in Question 1, how easy is the system to use?

- ☐ very easy
- ☐ easy
- ☐ adequate
- ☐ difficult
- ☐ very difficult

3. Operation of the system is easy to learn.

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

4. How would you rate the web user interface?

- ☐ very good
- ☐ good
- ☐ adequate
- ☐ poor
- ☐ very poor

5. How would you rate the viewing application?

- ☐ very good
- ☐ good
- ☐ adequate
- ☐ poor
- ☐ very poor

6. Is the region-of-interest tool (focus box) easy to use?

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

7. Do you feel that the region-of-interest selection capability is useful for reducing the download time of the image?

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

8. Using a dial up connection: I am happy at with the rate at which images are received.

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

9. Using a broadband connection: I am happy at with the rate at which images are received.

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

10. I am happy with the overall quality of images received.

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

11. From the received image and the region of interest selected, is the quality of image such that you are able to ascertain whether a lesion/wound/growth etc. is present on the subject's skin?

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

12. Following from question 11, are you able to distinguish between the different types of lesions/wounds/growths etc. ?

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

13. (Please ignore this question if you did not agree to question 11) Is the quality of image such that variations in colour/pigment of the lesion/wound/growth etc. are discernible?

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

14. (Please ignore this question if you did not agree to question 11) Is the quality of image such that variations in the surface of the lesion/wound/growth etc. are discernible?

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

15. (Please ignore this question if you did not agree to question 11) Is the quality of image such that variations in the shape/border of the lesion/wound/growth etc. are discernible?

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

Section 3 – Clinical Use and Benefits

16. Use of this system may improve the patient's overall health management.

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

17. Use of this system may improve the patient's wound and skin management.

- ☐ strongly agree
- ☐ agree
- ☐ no opinion
- ☐ disagree
- ☐ strongly disagree

18. Use of this system may reduce the amount of time needed to provide an accurate diagnosis.

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

19. I would you be happy to continue using the system on a regular basis

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

20. I am concerned about confidentiality and data security issues related to the use of the system.

☐ strongly agree
☐ agree
☐ no opinion
☐ disagree
☐ strongly disagree

Section 4 – Comments and Suggestions

21. How can the system be improved? Please give details of any changes that may make the system more useful.

....
....
....
....
....
....
....
....

22. Please comment on any difficulties using the system.

....
....
....
....
....
....

……………………………………………………………………………………
……………………………………………………………………………………

23. Please compare the performance of the system during downloads, between using (i) a modem connection, and (ii) a broadband connection.

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

24. Please feel free to leave any other comments (e.g. related to any of the questions asked in the previous sections).

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

Appendix D

UNSW HREC Application

ETHICS AND PRIVACY APPLICATION FORM FOR RESEARCH INVOLVING HUMANS

Please Note: Each question on this form has instructions and links to relevant documents and guidelines on how to answer that particular question as hidden text. To show the text with the hidden text effect, click symbol "¶" **(Show/Hide)** (situated next to the **"Zoom"** button) on the "**Standard**" toolbar. When hidden text is shown it is marked with a dotted underline. This text will not be seen on the printed version.

Please note the following:

1. This application must be completed electronically or typewritten
2. Complete all sections except those specifically not applicable
3. Use lay terms wherever possible
4. Do not alter the order of questions or layout of the application form
5. "Y" signifies Yes, "N" signifies No, and "N/A" signifies Not applicable
6. Some "Y"/"N" boxes have been reversed so take care in answering the questions
7. HREC refers to Human Research Ethics Committee

This form has been prepared in collaboration between Ms G Briody, Associate Professor M Grimm, Professor A Lloyd, Associate Professor J Watson and Ms M Wright of the Human Research Ethics Committees (HRECs) of the Universities of New South Wales and Sydney.

SECTION 1: ADMINISTRATION

This section is obligatory

1.1 (a) Full project title

A Novel Image Capture System for Use in Telehealth Applications

(b) Short name by which the project will be known

Image Capture System

(c) Name of Chief Investigator

Professor Branko Celler

(d) Provide a brief summary of the project in lay language (approximately 100 words)

A novel image capture and retrieval system has been developed for use in a range of telehealth applications in the home and in residential care facilities. The system is based around the JPEG2000 image standard and uses a standard camera protocol for image capture from any high resolution digital camera. When coupled with an image processing system such as that for pressure wound management, the system may provide a high level of clinical functionality suitable for a wide range of telemedicine applications in rural and remote sites.

(e) Outline the scientific merits of this study (including potential contributions to the body of knowledge and methodological rigor) (approximately 100 words)

This study will examine the feasibility of the image capture system in a residential care facility and test subjects. The main goals will be to gain understanding of the following: the impact of integrating the system with other existing healthcare modules; optimal conditions for image acquisition; feedback on system convenience, ease of use and patient and doctor satisfaction. The knowledge obtained will be used to improve on the current implementation to facilitate improved health outcomes.

1.2 Indicate the institutional ethics committee that you consider to be the primary one for this project. (In general, if the Chief Investigator is a University employee, then the University should be considered to be the primary site. If the Chief Investigator or participants are from a health care service, then the Area Health Service ethics committee should be considered as the primary site.)

UNSW HREC

1.3 (a) Has this project already been submitted to any other HREC(s)? [x] **N** [ ] **Y**

(b) Will this project be submitted to any other HREC(s)? [ ] **N** [x] **Y**

If you answered YES to (a) or (b), give the name of the HREC(s), and indicate the status of the application at each (i.e., submitted, approved, deferred or rejected).
Attach copies of the correspondence with each of the other HREC(s).
Please do not submit to more than one HREC <u>concurrently</u>.

Montefiore Homes Ethics Committee – not submitted until UNSW HREC approval granted.

1.4 List the following details of the Chief Investigator/Supervisor, any Co-Researcher(s), Associate Researcher(s) and Student(s).

Chief Investigator/Supervisor

Name	Branko Celler
Title	Professor
Qualifications	
Positions held: employed, conjoint/adjunct/visiting	Head, Biomedical Systems Laboratory, UNSW
Full mailing address (including building number)	
Telephone	
Fax	
E-mail	

Co-Researcher(s), Associate Researcher(s), Student(s) or other Personnel involved in the study (If appropriate indicate for each named person whether they are University staff, student or neither). If the named person is a student, nominate (in the Qualifications section) the degree for which he/she is enrolled.

Name	Ed Wyn Lim
Title	Mr.
Qualifications	Master of Engineering (Research)
Positions held: employed, conjoint/adjunct/visiting	Student
Full mailing address (including building number)	Biomedical Systems Laboratory, School of Electrical Engineering and Telecommunications, University of NSW, Sydney 2052, NSW
Telephone	
Fax	
E-mail	z2262590@student.unsw.edu.au

Name	
Title	
Qualifications	
Positions held: employed, conjoint/adjunct/visiting	
Full mailing address (including building number)	
Telephone	
Fax	
E-mail	

Name	
Title	
Qualifications	
Positions held: employed, conjoint/adjunct/visiting	
Full mailing address (including building number)	
Telephone	
Fax	
E-mail	

Name	
Title	
Qualifications	
Positions held: employed, conjoint/adjunct/visiting	
Full mailing address (including building number)	
Telephone	
Fax	
E-mail	

Insert additional boxes if necessary.

1.5 Who is the nominated Contact Person (from those listed in 1.4 above) for this protocol?

Name	Telephone Number	Email
Professor Branko Celler	9385 6595	b.celler@unsw.edu.au

1.6 Who is the person preparing this document?

Name	Telephone Number	Email
Prof. Branko Celler	9385 6595	b.celler@unsw.edu.au

1.7 In addition to the researchers named in 1.4 are there students involved as researchers in this project? ☒ N ☐ Y

If you answered YES, indicate the number of students covered by this study and the degrees which this study will contribute towards (i.e., Honours, Masters, PhD, etc.) **If the names are already known please include them.**

1.8 (a) Indicate the proposed date of commencement of the project.
Projects may not commence without the prior written approval of the HREC.

Date As soon as possible after HREC approval is obtained, no later than 1st August, 2006

(b) Indicate the proposed completion date of the project.

Date 27th October 2006

1.9 Indicate all location(s) at which the research will be undertaken.

Montefiore Nursing Homes at Maroubra and Hunters Hill

1.10 (a) Has this protocol received research funding/contracting or is this submission being made as part of an application for research funding/contracting? ☒ N ☐ Y

If you answered YES, list the funding/contracting bodies to which you have submitted, or intend to submit, this project. Attach a copy of the grant application(s), contract(s) or similar agreement(s).

Funding/Contracting body 1:
Funding/Contracting body 2:
Funding/Contracting body 3:

(b) What is the outcome of these funding/contracting application(s) (please tick the appropriate box)

Funding/Contracting body 1:	☐ **Approved**	☐ **Pending**	☐ **Refused**
Funding/Contracting body 2:	☐ **Approved**	☐ **Pending**	☐ **Refused**
Funding/Contracting body 3:	☐ **Approved**	☐ **Pending**	☐ **Refused**

(c) Will this study still be undertaken if funding is not successful? ☐ N ☐ Y

(d) If the title of the project submitted for funding is different from that listed under Q1.1(a), state it below.

Proceed to Section 2.

SECTION 2: NATURE OF RESEARCH
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 23-45)

This section is obligatory

2.1 The nature of this project is most appropriately described as research involving:- (more than one may apply):

	N	Y
– behavioural observation	☒	☐
– self-report questionnaire(s)	☒	☐
– interview(s)	☐	☒
– qualitative methodologies (e.g. focus groups)	☐	☒
– psychological experiments	☒	☐
– epidemiological studies	☒	☐
– data linkage studies	☒	☐
– psychiatric or clinical psychology studies	☒	☐
– human physiological investigation(s)	☒	☐
– biomechanical device(s)	☒	☐
– human tissue (see Section 11)	☒	☐
– human genetic analysis (see Section 11)	☒	☐
– a clinical trial of drug(s) or device(s) (see Section 12)	☒	☐
– Other (please specify in the box below)	☐	☒

Evaluation of a data collection modality and quality of images recorded

Proceed to Section 3.

SECTION 3: PARTICIPANTS AND RECRUITMENT
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 25-34)

This section is obligatory

3.1 (a) What is the age range of all participants involved in this study?

Elderly residents of Montefiore Residential Care Facilities, typically aged over 70

(b) If the participants include children (defined by statute for this purpose as anyone under 18) has a Prohibited Employment Declaration Form for the researchers ("criminal record check") been lodged with the University or hospital? (see http://www.kids.nsw.gov.au/check/) ☐ Y ☒ N

If you answered NO, give reasons why not.

This study focuses on care for elderly patients in a nursing home environment

3.2 Are the participants:-
(more than one may apply)

	N	Y
– in a teacher–student relationship with the researchers or their associates?	☒	☐
– in an employer–employee relationship with the researchers or their associates?	☒	☐
– in any other dependent relationship with the researchers or their associates?	☒	☐
– wards of the state?	☒	☐
– prisoners?	☒	☐
– refugees?	☒	☐
– members of the armed services?	☒	☐
– mentally ill?	☒	☐
– intellectually impaired?	☒	☐
– unconscious or critically ill patients?	☒	☐
– under the Guardianship Act 1987 (as amended)?	☒	☐
– in a doctor–patient relationship or a health giver–receiver relationship with the researchers or their associates?	☒	☐

If you answered YES to any of the above, provide details.

3.3 (a) What is the sample size for the study? Comment on how this sample size will allow the aims of the study to be achieved.

This is a qualitative study that seeks to evaluate the suitability of a data collection, storage and review format for digital images of clinical interest.

(b) How will the participants be recruited?

All residents in the Montefiore homes at Maroubra and Hunters Hill will be eligible to participate if they present with any lesions of clinical interest.

3.4 (a) Does recruitment involve a direct personal approach from the researchers to the potential participants? [x] N [ ] Y

If you answered YES, explain how the real, or perceived, coercion from researchers for potential participants to enrol has been addressed.

(b) Does recruitment involve the circulation/publication of an advertisement, circular, letter, etc? [x] N [ ] Y

If you answered YES, provide a copy and indicate where and how often it will be published.

3.5 Will participants receive any reimbursement of out-of-pocket expenses, or financial or other "rewards" as a result of participation? [x] N [ ] Y

If you answered YES, what is the amount or nature of the reward and the justification for this?

3.6 Is the research targeting any particular ethnic or community group?

[x] **N** [ ] **Y**

If you answered YES, which group is being targeted?

If you answered YES, is there an investigator who is a member of the Particular ethnic or community group?

[ ] **Y** [ ] **N**

If you answered YES to 3.6, has this project been planned in consultation with a representative of this group?

[ ] **Y** [ ] **N**

If you answered YES, who have you consulted and how do they represent this group?

If you answered NO, give reasons why you have not consulted.

Ethnicity is not relevant to this study, Digital images of clinical interest will be collected in de-identified form from any resident in the two nursing homes.

Proceed to Section 4.

SECTION 4: PRIVACY
Refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 52-53. For health related information refer to the Statutory Guidelines made under the *Health Records and Information Privacy (HRIP) Act 2002 (NSW) Statutory Guidelines on Research* via Privacy NSW HRIP Act and also the NHMRC overview document *The Regulation of Health Information Privacy in Australia* http://www.nhmrc.gov.au/publications/synopses/nh53syn.htm

This section is obligatory

4.1 Is there a requirement for the researchers to identify, collect, use, or disclose information of a personal nature *(either identifiable or potentially identifiable)* about individuals without their consent?

		N	Y
(a)	**from Commonwealth departments or agencies?**	☒	☐
(b)	**from State departments or agencies?**	☒	☐
(c)	**from other third parties, such as non-government organisations?**	☒	☐

If you answered YES to (a), (b) or (c), state what information will be sought and how many records will be accessed.

4.2 (a) Is there a requirement for the researchers to identify, collect, use, or disclose personal health information about individuals without their consent, which is identifiable or potentially identifiable? ☒ N ☐ Y

If you answered NO, you do not need to complete any more of Section 4. Go to Section 5

If you answered YES, indicate the reason(s)

- The project involves linkage of data ☐ Y
- Scientific deficiencies would result if de-identified information was used ☐ Y
- Other ☐ Y

Please provide details

4.3 Will the <u>health</u> information that is identifiable or potentially identifiable with respect to individuals be collected, used or disclosed without the consent of the individual(s) concerned? ☐ N ☐ Y

If you answered YES, indicate the reason(s)

- The size of the population involved in the research. ☐ Y
- The proportion of subjects who are likely to have moved or died since the health Information was originally collected. ☐ Y
- The risk of introducing bias into the research, affecting the generalisability and validity of the results. ☐ Y
- The risk of creating additional threats to privacy by having to link information in order to locate and contact subjects to seek their consent of the results. ☐ Y
- The risk of inflicting psychological, social or other harm by contacting subjects with particular conditions in certain circumstances. ☐ Y
- The difficulty of contacting individuals directly when there is no existing or continuing relationship between the organisation and the individuals. ☐ Y
- The difficulty of contacting individuals indirectly through public means, such as advertisement and notices. ☐ Y
- Other ☐ Y

Please provide details

4.4 Was this research the primary purpose of collecting the health information? ☐ Y ☐ N

If you answered YES, you do not need to complete any further questions in Section 4. Go to Section 5

If you answered NO, please provide details

4.5 Would the subjects have expected the researchers to use or disclose their health information for the purposes of this project? ☐ Y ☐ N

Please provide details

4.6 Explain why the collection, use or disclosure of this information is in the public interest, and why the public interest in the project substantially outweighs the public interest in the protection of privacy.

Proceed to Section 5.

SECTION 5: COLLECTION OF DATA AND DISSEMINATION OF RESULTS
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 52-53)

This section is obligatory

5.1 Will any part of the study involve recordings using audio tape, film/video, or other electronic medium ? N [ ] Y [x]
If you answered YES, what is the medium and how it will be used?

Digital cameras – for image acquisition and laptop computer for temporary storage of images. These will be transferred over a dial up facility to a secure server at the University of NSW

5.2 Does your research involve the secretive use of photographs, tape-recordings, or any other form of record-taking? N [x] Y [ ]
If you answered YES, provide details and a justification for the secrecy.

5.3 (a) How will the results of the study be disseminated (e.g. via publication in journals and presentations in scientific meetings)?

Graduate thesis, publication in journals

(b) How will feedback be made available to participants (e.g. via a newsletter)?

5.4 How will the confidentiality of the data, including the identity of participants, be ensured during collection and dissemination?

Patient images will be deidentified during collection and dissemination

5.5 Is there any possibility that information of a personal nature could be revealed to persons not directly connected with this research? N [x] Y [ ]
If you answered YES, provide details.

5.6 (a) What is the proposed storage location of, and access to, materials collected during the study (including files, audiotapes, questionnaires, videotapes, photographs)?

Computer Storage – in secure servers located at the Biomedical Systems Laboratory at the University of NSW

(b) Specify how long materials collected during the study (including files, audiotapes, questionnaires, videotapes, photographs) will be retained after the study, and how they will ultimately be disposed of.

Please ensure that the period of data retention stated here is appropriate to the nature of the proposed study. If the project involves clinical trial(s), the data should be kept for a minimum of 15 years (please refer to http://www.fda.gov/oc/ohrt/irbs/websites.html). If the projects do not involve clinical trial(s), the data should be kept for a minimum of 7 years after which time the data may be disposed of. *(Please also refer to National Statement on Ethical Conduct in Research Involving Humans, 12.11 for further requirements).*

Data is de-identified. It will be retained for the purpose of analysis and publication. However images may be used for additional research such as the development of algorithms to delineate wounds, identify gradations in texture and colour etc for improved tracking of longitudinal changes. There would appear to be no particular need to destroy what would be a useful databases of de-identified images that may prove useful for future research. If the HREC requires it however the images will be destroyed at the completion of this specific project.

Proceed to Section 6.

SECTION 6: RISKS AND BENEFITS
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 51)

This section is obligatory

6.1 (a) Could participation in the research adversely affect the participants? [x] N [ ] Y

If you answered YES, complete 6.1 (b) and 6.1 (c). If you answered NO go to 6.2

(b) Could the research induce any psychological distress in the participants? [x] N [ ] Y

(c) Could the research cause any physical harm to the participants? (e.g. from physically invasive procedures or from drug administration, etc) [x] N [ ] Y

If you answered YES to (b) or (c) describe the aspect(s) of the research and all the risks involved. Indicate the rate at which these risks are expected to occur. Indicate what facilities and trained personnel are available to deal with such psychological or physical problems.

6.2 Will the true purpose of the research be concealed from the participants? [x] N [ ] Y

If you answered YES, outline the rationale and provide details for the concealment. Provide details of the debriefing. (If you do not intend to debrief, give reasons why not).

6.3 Are you doing research on patients (i.e. subjects receiving health care)? [ ] N [x] Y

If you answered YES, list the procedures/techniques which would not form part of routine clinical management.

Subjects are residents in Residential Care Facilities and are subject to normal care. The recording of digital images and their review by the patients own GP or another authorized member of the clinical care team represents a potential improvement on existing care by permitting a clinician to remotely view the images and advise on treatment.

6.4 Is this research expected to benefit the participants directly or indirectly? [ ] N [x] Y

If you answered YES, provide details.

Outcome of research is expected to yield an improved visual diagnostic method (image capture via consumer digital cameras), which will be incorporated into existing telemedicine system developed by the researchers for use in Residential Care Facilities or rural and remote community heath centres.

Proceed to Section 7.

SECTION 7: PARTICIPANT INFORMATION AND CONSENT
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p.12-13, p.28-29, p. 40-42, p.44-45, p.47-50, p.54)

This section is obligatory

7.1 Will a Participant Information Statement be provided? [x] Y [ ] N

7.2 Will written consent be obtained? [x] Y [ ] N

If you answered NO to either 7.1 or 7.2, give reasons why not.

7.3 In the case of participants who may not be fluent in English or who have difficulty understanding English, will arrangements be made to ensure comprehension of the Participant Information Statement and Consent Form? [ ] Y [x] N

If you answered NO, give reasons. If you answered YES, what arrangements have been made?

All residents who will be asked to participate in this study are fluent in English

7.4 (a) Do the Participant Information Statement and Consent Form have:-

- **the first page of the Participant Information Statement and Consent Form printed on appropriate institutional letterhead?** [x] Y [ ] N
- **the title of the project on every page, including the Revocation of Consent? (if one is required) (Use a short title as appropriate)** [x] Y [ ] N
- **the page numbers expressed as page 1 of .., 2 of .., 3 of .. etc?** [x] Y [ ] N
- **an assurance that participation is voluntary and participants are permitted to withdraw from the project at any time without penalty?** [x] Y [ ] N
- **the name and telephone number of an appropriate researcher?** [x] Y [ ] N
- **a telephone number, fax number and E-mail address for the HREC, should a participant wish to make a complaint about the conduct of the research project?** [x] Y [ ] N

(b) How has the possibility of withdrawal from the study been addressed in the Participant Information Statement and Consent Form?

Yes

Proceed to Section 8.

SECTION 8: CONFLICT OF INTEREST AND OTHER ETHICAL ISSUES
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 51–54, Appendix 2)

This section is obligatory

8.1 Are any "conflict of interest" issues likely to arise in relation to this research? ☐ N ☒ Y

If you answered YES, provide details.

There is an existing Research Management Agreement between MedCare Systems Pty Ltd and the University of NSW for possible commercial exploitation of research undertaken at the Biomedical Systems Laboratory. This project however, at this stage is strictly a thesis project for an ME (Research) student, Mr. Ed Wyn Lim

8.2 Do the researchers have any affiliation with, or financial involvement in, any organisation or entity with direct or indirect interests in the subject matter or materials of this research? ☒ N ☐ Y

(Note that such benefits must be declared in the Participant Information Statement.)
If you answered YES, provide details.

8.3 Do the researchers expect to obtain any direct or indirect financial or other benefits from conducting this research? ☒ N ☐ Y

(Note that such benefits must be declared in the Participant Information Statement.)
If you answered YES, provide details.

8.4 (a) Have conditions already been imposed upon the use (eg. publication), or ownership of the results (eg. scientific presentations) or materials (eg. audio-recordings), by any party other than the listed researchers? ☒ N ☐ Y

(b) Are such conditions likely to be imposed in the future? ☒ N ☐ Y

If you answered YES to (a) or (b), provide details.

Proceed to Section 9.

SECTION 9: DESCRIPTION OF PROJECT
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 13)

This section is obligatory

9.1 Describe the project using lay terms wherever possible, including the aims, hypotheses, research plan and potential significance. Where relevant, provide the projected number, sex, and age range of participants (including inclusion/exclusion criteria). You must satisfy the HREC that the study is scientifically valid and conducted in accordance with the accepted ethical principles governing research involving humans.
The description must be no longer than 2 pages and must be in a font size of at least 10 points.

A Novel Image Capture System for Use in Telehealth Applications
A novel image capture and retrieval system has been developed for use in a range of telehealth applications in the home and in residential care facilities. The system is based around the JPEG 2000 standard and uses the PTP protocol for image capture from any high resolution digital camera and the Kakadu suite of JPEG2000 utilities to serve the collected images via a proxy server over any available communication channel from telephone lines to broadband services.

When coupled with an image processing system such as the AMWIS system for pressure wound management, the system provides a high level of clinical functionality suitable for a wide range of telemedicine applications in rural and remote sites.

This project is concerned with the trialing of a low-cost imaging system, able to capture an image with a consumer digital camera, and delivering the image efficiently over the Internet to the medical practitioner in charge of the patient. The emphasis is on the acquisition and delivery of images from patient to practitioner. The system is thus potentially a good complement to a dedicated wound imaging system like the AMWIS system being employed in several sites around Australia.

Technical details are described in an attached paper which has been accepted for presentation at the IEEE Engineering in Medicine and Biology Conference in Ney York later this year.

9.1 (continued)

Proceed to section 10.

SECTION 10: FIELD-BASED RESEARCH (i.e., CONDUCTED OFF CAMPUS OR OUTSIDE A HEALTH SERVICE) INCLUDING RESEARCH CONDUCTED OUTSIDE AUSTRALIA
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p.14, p.31-32)

This section must be completed for all applications involving EITHER field-based research OR research to be carried out in countries outside Australia (eg. in a school, a corporation, a government department an Aboriginal and Torres Strait Islander community or research in a another country).

10.1 Does this section apply to your research? ☐ N ☒ Y

If NO, Go to Section 11

10.2 Have you obtained formal permission from relevant authorities for entry to the area to carry out research (e. g., national or local government bodies, organisations of local communities)? ☐ Y ☒ N

If you answered YES, name the relevant authorities and attach the relevant correspondence.

If you answered NO, give reasons.

Discussion have been had with the Montefiore Homes and they have indicated that they are happy to participate. Once HREC approval is received, they will present the study and the HREC approval for consideration by their own institutional ethics committee

10.3 If research is proposed among members of specific organisations, have you sought approval from those organisations (e. g., church groups, national associations, etc)? ☒ Y ☐ N

If you answered YES, name the relevant authorities and attach the relevant correspondence or letter of support.

Montefiore Homes. Discussions and expressions of interest in participation have taken place with senior staff at Montefiore Homes. No exchange of letters as yet as these are pending on HREC approval.

If you answered NO, give reasons.

10.4 Does the research involve individuals or groups of people who are not formally organised (e.g., people living in a village or town, etc)? ☒ N ☐ Y

If you answered YES, indicate the context of the research. How will you obtain access to participants? Indicate any ethical issues that you can foresee in this approach.

10.5 Will your research necessarily involve the acquisition of objects of valuable cultural property (e. g., carvings, paintings, etc)?

[x] N [ ] Y

If you answered YES, give details of arrangements with owners of the property with regard to access to/acquisition of these items, where appropriate.

10.6 Will your research necessarily involve any activities that are likely to be seen by research participants and/or members of their local communities as in conflict with local practices and customs (e.g. regarding religious or ritual participation)?

[x] N [ ] Y

If you answered YES, provide details.

Proceed to Section 11.

SECTION 11: RESEARCH INVOLVING BLOOD, TISSUE, ETC.
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p.33, p.43-50)

This section must be completed for all research involving blood or tissue samples, or involving physical hazards.

11.1 Does this section apply to your research? [x] N [ ] Y

If NO, Go to Section 12

11.2 Will human blood or tissue be used in the research? [ ] N [ ] Y

If you answered YES, what procedures are in place to minimise the infectious and other risks to participants and researchers?

11.3 Will human embryos, fetal tissue, or placental tissue be involved? [ ] N [ ] Y

If you answered YES, provide details.

11.4 Has this blood or tissue already been collected and stored? [ ] N [ ] Y

If you answered YES, what was the original purpose of collection for the stored blood or tissue you seek to use?

11.5 Describe the proposed storage arrangements of the blood and/or tissue samples collected. Indicate how long the blood or tissue will be kept. Indicate how the samples will be disposed of upon the completion of the research.

11.6 Will genetically modified organisms or other gene modification techniques be used in the research? [ ] N [ ] Y

If you answered YES, provide details. Describe the procedures, which are in place to minimise the risks to participants and researchers.

11.7 Will toxins, mutagens, teratogens or carcinogens be used? ☐ N ☐ Y

If you answered YES, provide details. Describe the procedures, which are in place to minimise the risks to participants and researchers.

11.8 Will biohazardous material be used? ☐ N ☐ Y

If you answered YES, provide details. Describe the procedures, which are in place to minimise the risks to participants and researchers.

11.9 Will participants or researchers be exposed to ionising radiation? ☐ N ☐ Y

If you answered YES, provide details of the radiation exposure, including a quantitative assessment of the absorbed dose, supported either by dosimetric calculations or by other information. Describe the procedures, which are in place to minimise the risks to participants and researchers. The study should also be approved by the relevant institutional Radiiation Safety authority.

Proceed to Section 12.

SECTION 12: CLINICAL TRIALS OF DRUGS OR DEVICES
(refer to the National Statement on Ethical Conduct in Research Involving Humans, p. 35-38, and also to Therapeutic Goods Administration, http://www.tga.gov.au)

This section must be completed for all applications involving clinical trial(s).

12.1 Does this section apply to your research? [X] **N** [ ] **Y**

If NO, Go to Section 13

12.2 (i) Is the research being conducted under the Clinical Trial Notification Scheme (CTN)? [ ] **N** [ ] **Y**

(ii) Is the research being conducted under the Clinical Trial Exemption Scheme (CTX)? [ ] **N** [ ] **Y**

(iii) Is the research using only approved drug(s)/device(s) in accordance with Therapeutic Goods Administration Approved Product Information? [ ] **Y** [ ] **N**
(Note reversed order of the responses)

12.3 (a) Will this research be undertaken on behalf of (or at the request of) a pharmaceutical company, or other commercial entity, or any other sponsor? [ ] **N** [ ] **Y**

If you answered YES, provide details of the name of the sponsor (and co-sponsors if any) ?
This information should be included in the Participant Information Statement and Consent Form.

Will the sponsor(s) provide any support in money or kind?
Provide details.

(b) If you answered YES to (a) will that entity undertake in writing to abide by either the Medicines Australia Guidelines for Injury Resulting from Participation in an Industry-Sponsored Clinical Trial (www.medicinesaustralia.com.au) or the ABPI Clinical Trial Compensation Guidelines? [ ] **Y** [ ] **N**

If you answered NO to this question, provide details.

(c) If you answered YES to (a), will that entity undertake in writing to indemnify the institution, the HREC(s) and the researchers ? [ ] **Y** [ ] **N**
(If you answered YES, a copy of the appropriate deed or letter of indemnity should be included with the application).

If you answered NO to this question, provide details.

(d) **If you answered YES to (a), (b) or (c), does the sponsor hold a current insurance policy to cover this project?** ☐ Y ☐ N
(*If you answered YES, provide a certificate of currency*).

If you answered NO to this question, provide details.

12.4 List any drugs or devices to be used, and their TGA approval status both in Australia and overseas ☒ NA

12.5 How many participants are projected to be enrolled into the trial at this site and in total? (Please give a single figure for each, not a range)

12.6 What is the projected duration of the trial, from first enrolment to the last protocol interaction with the last enrolled subject (in years)?

12.7 If all projected participants complete the protocol:

(a) what total payment will be received from the sponsoring company?
(Please give a single figure, not a range)

(b) what additional "in kind" support (ie free drug, equipment, etc), if any, will be provided by the sponsoring company?

For instructions on how to obtain TGA approval, please refer to **http://www.tga.gov.au**.

Proceed to the Section 13.

SECTION 13. DECLARATION OF RESEARCHERS

I/we apply for approval to conduct the research. If approval is granted, it will be undertaken in accordance with this application and other relevant laws, regulations and guidelines.

Signature of Chief Investigator or Supervisor

Name: Professor Branko Celler (print)

Signature:

Date: 4th July 2006

Signature of Associate Researcher(s) or Student(s)

Name Mr. Ed Wyn Lim. (print)

Signature:

Date: 4th July 2006 ....

Name .... (print)

Signature: ....

Date: ....

Name .... (print)

Signature: ....

Date: ....

Signature of appropriate senior officer NOT ASSOCIATED with the research (e.g. Head of School/ Department/Unit/Dean of Faculty or Head of Division).

After careful consideration and appropriate consultation, I have reviewed the attached HREC application, including the Participant Information Statement and Consent Form. I am satisfied that the scientific merit of this work justifies its being performed and that the information which will be obtained justifies the inconvenience and risks to participants.

Name: Tim Hesketh ....(print)

Title: Associate Professor .... (print)

Position: Head of School... (print)

Signature:

Date: 4th July 2006.

CHECKLIST FOR FULL ETHICS APPLICATION

The following documents are to be attached as indicated in the Guide to Applicants.

Check N/A if not applicable.

Have you included the **original copies (plus the number of copies required by your HREC)** of the following:

	Y	N/A
Original application	[x]	
Consent form(s)	[x]	[ ]
Participant Information Statement (s)	[x]	[ ]
Recruitment advertisement/circular	[ ]	[x]
Evidence of permission to conduct research in other locations	[ ]	[x]
Evidence of approval/rejection by other HREC(s), including comments and requested alterations to the protocol	[ ]	[x]
Copy of questionnaire(s), survey questions, interview topics to be covered etc.	[ ]	[x]
Statement from a medical/paramedical practitioner accepting responsibility for specific procedures.	[ ]	[x]
Risk management unit report regarding genetically-modified organisms, biohazards, Ionizing radiation, lasers or carcinogens	[ ]	[x]
One copy of the grant application with appropriate clearance forms as requested by the Research Office **(Refer to your local requirements)**	[ ]	[x]
Any form requiring signature by the HREC (one copy) e.g. CTN/CTX Forms	[ ]	[x]