

Copyright

by

Eoin Meehan

2006

Are you looking at me?

Eoin Meehan

A thesis submitted to the University of Dublin, Trinity College
in partial fulfillment of the requirements for the degree of
Master of Science

May 2006

Are you looking at me?

**Approved by
Supervising Committee:**

Declaration

I, the undersigned, declare that this work has not previously been submitted to this or any other University, and that unless otherwise stated, it is entirely my own work.

Eoin Meehan

Dated: May 28, 2006

Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this thesis upon request.

Eoin Meehan

Dated: May 28, 2006

Acknowledgements

I would like to acknowledge Dr. Vincent Weldon of the Physics Department, Trinity College Dublin, for kindly responding to my initial queries about using lasers and setting me in the right direction. I would like to thank Dr Louise Bradley of the Physics Department, for ensuring I and my project operated safely and for giving me laboratory space and much support in building the prototype. I would like to thank Gillian O'Neill also of the Physics Department for helping me source optical equipment and making sense of optical terminology. I would like to thank John D. Maniraj for giving me 8051 sample code on which I was able to base the SPI communications module. I would like to thank my supervisor Mr James Garland and my second reader Dr Stefan Weber for their support and encouragement. I would like to thank my colleagues in the Ubiquitous Computing class for their general support and good humour during the course. Finally, my thanks to Meaux for her faith and support.

Eoin Meehan

University of Dublin, Trinity College

May 2006

Abstract

The aim of Ubiquitous Computing is to embed devices into the fabric of everyday things and have these devices seamlessly interact with users to enhance their experience of the world. However, this means these devices cannot use the traditional user interfaces such as keyboard, mouse and screen. Devices must be able to infer what the user wants or how to assist the user by deducing the context of the interaction even when the user might not even be aware the device exists.

Context is a wide term and in ubiquitous computing usually means determining the users physical and temporal position as well as proximity to interesting things and events. Probability distributions can also be used to infer context based on previous user actions and movements.

This dissertation proposes that as humans use eye-gaze as a control mechanism in conversations, eye-gaze can be also used in interaction with ubiquitous computing devices. Simply looking at the device the user wishes to address produces an important item of context. All other devices can ignore the impending command.

Detecting eye-gaze is a rich research area, but this dissertation proposes to use a scanning laser to detect eye gaze, rather than video image processing. This has a number of advantages including cost and miniaturisation potential. If devices are to become truly ubiquitous, then the devices have to be cheap enough to embed in everything without regard to cost.

Contents

Acknowledgements	vi
Abstract	vi
List of Tables	xi
List of Figures	xii
Chapter 1 Introduction	1
1.1 The Problem in Context	1
1.2 The Problem of Context	2
1.3 Conversation	4
1.4 Ideal UbiCom Interface	6
1.5 The Illusion of Intelligence	6
1.6 Objectives	7
Chapter 2 Background	9
2.1 State of the Art	10
2.1.1 Video Image Processing Based Eye-Gaze Detection	10
2.1.2 Laser Based Eye-Gaze Detection	13
2.2 Related Work	15

Chapter 3 Design	18
3.1 Introduction	18
3.2 Design for Test	20
3.3 Laser technology	21
3.4 Laser Safety	23
3.5 Red-eye and retroreflectors	26
3.6 Scanning	28
3.7 Light detector	29
3.8 Analogue to Digital converter	29
3.9 Processor	29
Chapter 4 Implementation	30
4.1 Introduction	30
4.2 Hardware	31
4.2.1 Processor - RITA51J Evaluation Board	31
4.2.2 Analog to Digital Converter - AD7718	33
4.2.3 Rotating Polygon Laser Scanner	34
4.2.4 Light Detector	36
4.2.5 Lasers	36
4.3 Software	37
4.3.1 Supplied Software	37
4.3.2 Written Software	38
4.4 Building the prototype	39
4.4.1 Interconnecting the boards	41
4.4.2 AD Values in Detail	44
4.4.3 Timings	45
4.4.4 Loss of Power in the Returned Signal	45

Chapter 5 Evaluation	47
5.1 Design	47
5.1.1 Cost	47
5.1.2 Power	48
5.2 Implementation	48
5.3 Original Objectives	51
5.4 Security Considerations	53
Chapter 6 Conclusions	57
6.1 Summary	57
6.2 Further work	57
6.2.1 Optimising the Prototype	58
6.2.2 Towards Commercialisation	58
6.2.3 Disabled and Blind Assistance	59
6.3 Conclusion	60
Bibliography	61

List of Tables

3.1	Laser Classes	25
4.1	AD values returned	44

List of Figures

2.1	RBS system overview	15
2.2	Eye and reflected glints	16
3.1	Original design overview	19
3.2	Basic Laser Operation (Image Copyright 1994-2006, Samuel M. Goldwasser)	22
3.3	Retroreflector Principle	27
3.4	Raster Scan	28
4.1	Overview of the prototype	31
4.2	Rita51J evaluation board	35
4.3	AD7718 evaluation boards	35
4.4	Rotating Polygon Scanner	37
4.5	Two reflector targets	40
4.6	Original SPI interconnect	41
4.7	Final connections	42
4.8	SPI Data Register Read cycle	43
4.9	Scanner timings	46
5.1	Simplified diagram of the eye	51

Chapter 1

Introduction

1.1 The Problem in Context

In the early days of the computer era (c. 1950-1960) the interfaces to computer systems were designed to facilitate the system rather than the user. Data had to be coded in a convenient medium for the system to understand by specially trained operators and early examples include punch-card and paper-tape [19]. This medium was then submitted to the system for processing. This processing was usually “batched” into overnight runs. The output from the system was either more punch-cards or print-outs. This didn’t matter at the time as the systems never had to communicate directly with the end-user. Tasks were given to trained operators who would prepare the data for input, submit the data, monitor the run and collect the output.

The next stage of interface development was the introduction of the Visual Display Unit (VDU). This afforded the operators more direct control of their systems and we begin to see operators interacting with the systems to modify, control and manage tasks as they are running. However, special knowledge was still needed to communicate with the system as the interface was a simple command line.

From 1980 onwards, we begin to see the deployment of the WIMPS interface (Windows, Icon, Mouse, Pointing System) [34]. This was an interface designed to allow users interact directly with the system without the intervention of skilled intermediaries (operators) or specialised knowledge. It has become the de facto interface to most systems with commercial products from Microsoft, Apple, Sun, etc.

In Ubiquitous Computing [39] the next stage of development of the user interface is here; the hidden user interface. Now it is possible for users to interact with computer systems without having to learn specialised techniques and even, perhaps, without realising they are interacting. These computer systems will be built into the fabric of the world and people's primary experience of computation will be with a range of computationally-enhanced devices [12].

For the purposes of this paper, the term "ubicom" will mean "ubiquitous computing".

1.2 The Problem of Context

Much work has been done in investigating different ways a user can interact with a ubicom system by making that system aware of context, for example Chen and Kotz [8]. Traditional systems have interacted with users using explicit commands and implicit context. For instance, a user double-clicking on a image instructs the system to display that image to the user; the implied context is that the user wants to see the image now, because he is sitting in front of the screen and has issued the command on the keyboard or mouse.

Ubicom systems have to contend with the possibility that the user may not explicitly be instructing the system and that the instruction may mean different things depending on the context.

However, what is context? Context can have many meanings. Traditionally in computer science, context has been used in terms such as "context sensitive help", "contextual search", "multitasking context switch", "psychological contextual perception", and so on.

Some researchers have tried to define context more accurately by creating categories such as *computing context*, *user context*, *physical context* and *time context*. Chen and Kotz [8] give a useful definition of context as:

Context is the set of environmental states and settings that either determines an application's behavior or in which an application event occurs and is interesting to the user.

They further suggest that context aware computing falls into two categories:

Active context awareness: an application automatically adapts to discovered context, by changing the application's behavior.

Passive context awareness: an application presents the new or updated context to an interested user or makes the context persistent for the user to retrieve later.

However, Chen and Kotz were surveying existing context aware applications, especially mobile application. They were really focusing on what was out there, rather than attempting to define what context means. Consequently many applications limit what they mean as context to the users current and past physical and temporal locations, proximity to objects and events and probability distributions based on previous actions and events.

Dourish [12] discusses the problem of defining context with relation to ubiquitous computing. He argues that most current work is based on presenting context as a representational problem. That is, one that can be encoded and represented. In this view, context is a form of information; It is delineable and stable; Context and activity are separable - activity happens "within" a context.

Dourish proposes that context is an interactional problem. In this view, context is a relational property; Its scope is defined dynamically; It is an occasioned property, unique to each occasion, activity or action; Context arises from the activity, actively produced, maintained and enacted in the course of the activity at hand.

Dourish reports on the work of Harvey Sacks, founder of the field of Conversational Analysis [14], a subfield of ethnomethodology. Sacks found that people are constantly working at defining what is ordinary and what is remarkable in conversation [14]. Dourish suggests that context is also something that is constantly being defined, something that is an occasioned property of action. Like ordinariness, context is managed moment by moment, achieved by those carrying out some activity together, and relative to that activity and to the forms of action and engagement that it entails.

1.3 Conversation

One way of treating the interaction of a user with ubicom devices is to look at the interaction as a multi-party conversation. In conversation, humans use eye-contact and gaze-length to assist in controlling the interaction [37]. In multi-party conversations, eye-contact is much higher than in dyadic (two-party) conversations and this is believed to be because the natural “turn-taking” of dyadic conversations is missing. The speaker must “hold” the floor until he is ready to handover to the next speaker - gaze also controls who is the next speaker [38]. In effect, when people are speaking in multi-party conversations they look at the listener. The listener is also looking at the speaker, and the other participants, trying to identify cues.

Looking at a multi-party conversation in progress, it's obvious there are rules being obeyed and processes being followed. These rules not only apply to the mechanics of “turn-taking” or interruptions, but also to the “parsing” and interpretation of what is being said. An important aspect of this is that we understand the concept of “ordinariness” as mentioned above [12]. During a conversation, calculating what is “ordinary” is an on-going, mutual process performed by the participants. It allows the participants to constantly parse the conversation stream, and tacitly understand elements of the conversation. It allows the participants focus on what is “remarkable” and make that more important.

The author of this dissertation proposes that in a user's interaction with ubicom systems,

this interaction is, in effect, a multi-party conversation. The user is talking to the systems, and the systems may be talking to each other. The interaction process should abide by the rules of conversation and “be ordinary”. For example, if a customer in a restaurant wants to order coffee, he catches the eye of a waiter, when the waiter comes toward him he says “two coffees, please”. The waiter nods and heads off to the kitchen. The “ordinariness” of this situation is that it normal for the customer not to have to introduce himself, identify that he is a customer, request the coffee and to bring milk and sugar, indicate that he has some means of paying for it, and tell the waiter where he wants that coffee, ideally on the table and in a cup. There is also a temporary dispensing of the formalities of etiquette. But this interaction would not have taken place if the customer had not looked at the waiter. Looking at the front-door and shouting “two coffees, please” would probably just have annoyed the other customers and confused the staff.

So humans have adopted the use of eye-gaze as a technique to manage multi-party conversations. Eye-gaze has the added benefit that it does not interfere with the audio “channel” and can run in parallel. It has also become “ordinary” in that we use it without thinking.

The author’s postgraduate diploma project [25] (jointly with fellow student Joe McKnight) involved building a gesture-recognition system to interpret basic gestures that would allow a user control a domestic TV. They discovered that the gesture recognition worked fine when they had a single, dedicated “target”. However, if they expanded this system to control multiple targets, how could the user tell the system which target the gesture is meant for? They would have to develop a set of special gestures to tell the system which target they were now controlling.

The author of this dissertation proposes that it is possible to use eye-gaze in interacting with ubicom devices. Whatever means employed to control the device, whether it be voice-commands, remote control, or other technique, the user can begin by first looking at the device he wishes to address. This is a natural action and operates in conjunction with the control technique. The user cannot look at more than one device at a time, hence this action

is unambiguous.

1.4 Ideal Ubicom Interface

In the discussions above, it can be seen that ubicom devices may not use traditional

1. It must be simple for the user to use (minimum or ideally no learning required)
2. It must be ordinary for the user to use (unremarkable)
3. It must be natural for the user to use (comfortable both physically and intellectually)
4. The user must not have to wear anything special (e.g. special interface devices)
5. The user must not have to do anything out of the ordinary (e.g. special gestures)

This project aims to capture one aspect of normal human interaction - that of eye-gaze and in such a way as to fulfill the characteristics mentioned above.

1.5 The Illusion of Intelligence

In designing applications for ubiquitous computing, it is often tempting to choose powerful hardware platforms and involved software components. Indeed, there are applications where this is necessary. However, it is also possible that many devices, doing very simple things can combine to give the “illusion of intelligence”. For instance, it’s possible to build a coffee machine that detects a sleeping person is waking up in the morning, and quietly asks him “Coffee?”. He answers yes, and the machine merrily makes his morning coffee. To achieve this level of seeming intelligence, the coffee machine only requires a motion sensor, a timer, a digital recording of “Coffee?” and the ability to pick a high treble note out of the voice response (the high treble in the word “yes”). Yet to the outside observer, it looks as if the machine is intelligent! This concept was discussed as far back as 1950 by the eminent

Alan Turing [36] where he discusses the question “Can machines think?”. If a system can simulate intelligence (or seem intelligent to a user) does it really matter that it is, in fact, not intelligent? The author believes this is an important issue for ubicom applications as it can lead to the building of more cost-effective, robust and long-lived components for the ubiquitous computing environment.

1.6 Objectives

The authors main objective of this project is to see if it is possible for a device to recognise that a user is looking directly at it, using a positive aspect of the photographic red-eye effect. This is something that can be seen in family photographs, where the subject in the photograph acquires glowing red eyes. This effect is caused by light from a camera flash unit, entering the eye, bouncing off the retina, and returning along a similar path back to the source. In the same manner, using an eye-safe laser, it should be possible to emit a beam and detect it returning from the retina to the emitter.

Ideally, the laser should be infra-red so that the scan is invisible and therefore unobtrusive, but this raises issues of laser safety as will be discussed.

For the purposes of this project, a visible-light red laser will be used along with reflectors and will attempt to simulate gaze detection and proof-of-concept.

The following is a list of the objectives for this project which will be discussed in the evaluation section:

1. Does a spherical retroreflector simulate a human eye successfully?
2. Can the reflected laser energy from the retroreflector be detected by a photoreceptor, separately from the normal energy reflected from a surface?
3. Can the reflection be detected continuously enough for the processor to realise it is a gaze?

4. What is the minimum laser power required for this simulation to work?
5. What type of photoreceptor works best?
6. What is the best method to cause the laser to scan? A raster-like pattern or a diffused “fan” pattern?
7. Calculate optimum “scan” rates
8. Determine if energy required can be brought within eye-safe limits for a visible light laser.
9. Determine if this limitation can be implemented using an infra-red laser (infra-red lasers do not invoke the human blink-response)

Chapter 2

Background

Visual input is very important to humans so we tend to look at whatever it is that has our attention. We focus our eyes on it and look at it for some length of time; this is what we mean by gaze. Even in our vocabulary, to “focus on” means to “give all our attention to”. Our eyes are designed so that whatever we are focusing on is in sharp detail, and our peripheral vision is blurred.

The term “gaze-tracking” or “gaze-location” is used to describe capturing where a person is looking. Being able to track and capture a users gaze has a number of useful applications including supporting human-computer interaction (HCI) and supporting applications for the infirm or physically disabled where a user might only have control of their eyes.

Most of the research appears to be *user-centric*; that is, it is concerned with *where* the user is looking and tracking the gaze of the user. The research is divided into two groups. The first and most popular group uses video image processing and the second uses laser technology. Video image processing involves both bitmap image manipulation and mathematical processing to infer gaze location. Laser processing tends to involve more physical manipulation of the incident beam and interpretation of the reflected beam.

2.1 State of the Art

2.1.1 Video Image Processing Based Eye-Gaze Detection

Matsumoto et al, [24] designed a system that incorporates a stereo camera pair into a computer monitor to measure 3D position of the facial features of the user. The measurement system has three modules:

- Face Finding Module
- Face Measurement Module
- Facial Behaviour Recognition Module

The Face Finding Module extracts a face region from a stereo image pair using a standard edge template of a face. The facial features are then extracted. The Face Measurement Module extracts the various facial information such as head position, gaze direction, blinks etc. It calculates gaze direction by firstly calculating the centre of the eyeball, then calculating the radius of the eyeball, then detecting the iris, then deducing the gaze from the offset position of the iris from the center of the eye. The Face Measurement Module extends the ability of the system further by recognising such gestures as winking and talking.

Takemura et al [33] took this system and designed a real-world application for it. The concept is “what you look at is what you control”. The system uses gaze to select a target of the operator, which could, for example be a table lamp. The system is limited to the user sitting at a desk in front of a computer monitor on which the cameras are mounted. The user looks at the target and then clicks the mouse to affect the target.

The system requires that the location of all targets be defined in advance and the area of those targets is a sphere. The system calculates when the gaze vector passes through the sphere of the target. The system operates very successfully calculating the gaze vector to within 5 degrees and the whole process takes around 10ms. One successful experiment

was using one keyboard and mouse with two different monitors on two different computer systems. Obviously the one being looked at is the one the user wanted to interact with.

The disadvantage of this system is that it uses high processing power (a 2 Ghz Pentium 4) and the target device locations have to be predefined. If the targets are moved the system will not work. It also requires the use of a stereo camera pair.

Heinzmann and Zelinsky [15] designed a 3D facial pose and gaze point estimation using a monocular camera. They argue that other 3D based systems cannot recover from tracking errors where for example, the subjects head moves too quickly. Their system calculates the pose of the subjects head, then the location of the iris and the inner and outer corner of the eyes are tracked. The pose is calculated by finding facial features in the 2D image and translating them into a 3D model.

This paper is interesting because it shows the progression from stereo camera processing to monocular camera processing with a reported increase in effectiveness. However, this system still requires high processing capability and is based on the MEP vision tracking system produced by Fujitsu.

Magee et al [22] designed a system that runs on an average PC using an inexpensive USB camera. Their motivation was to design a system that could be used by people with severe physical disabilities, but without specialised expensive equipment or active infra-red illumination. Furthermore, the system had to be able to run on a “home-grade” PC in real-time. This paper also shows a further progression from using a relatively expensive monocular video camera to an inexpensive webcam.

The system is limited to detecting whether the user is looking left, right or straight ahead but this is an important example given the idea of building the Illusion of Intelligence as mentioned in chapter 1.

Their system has two components:

1. Face Tracker

2. Eye Analysis

The Face tracker uses colour and motion analysis along with template matching to find the location of the face in the video image. The Eye Analysis module reports on whether the user is looking left, right or straight ahead. To do this it takes the left-eye image, mirrors it and subtracts from the right eye image. The signed difference between the two images shows distinct pixel areas where the pupils are in different locations in each image. The signed differences are projected onto the x-axis. A strong positive peak followed by strong negative peak indicates left direction, while a strong negative peak followed by a strong positive peak indicates right direction.

As part of this paper, a simple game “BlockEscape” was constructed in which the user moves a block left or right on a screen. The system was also tested with a user with cerebral palsy where the user had to move a window left and right across a screen and this was successful.

Another technique used by researchers is to use active infra-red (IR) illumination in conjunction with video image processing. Zhu and Ji [40] designed a system using active IR illumination that maps user gaze position to screen coordinates to improve the user interaction with a graphic display.

Their system uses two IR illuminators - two rings of IR LEDs. One ring is located on the optical axis and the second off the axis (i.e. an inner ring of LEDs and an outer ring of LEDs). The rings are turned on and off alternatively and this allows the system to detect both the pupil of the eye and the glint (reflection off the cornea). The system then adds this information to a face pose detector and applies the results to a General Regression Neural Network. The result returned is the section of the screen that user is looking at. Using IR illumination allows the system to use narrow optical bandpass filters to eliminate most normal ambient visual light and so improves accuracy and reduces processing.

Park [29] also uses IR illumination and glint detection but with two cameras - a wide field camera and a narrow-field camera. The wide field camera is used to detect the facial features and this information is used to maneuver the narrow-field camera to focus on the eyes using pan and tilt. This system has the advantage that no calibration is needed, but does rely on a two-camera system, with the second camera having pan-tilt control.

2.1.2 Laser Based Eye-Gaze Detection

Irie et al [18] designed a laser-based eye-tracking system which tracks the horizontal, vertical and torsional movement of one eye. It does this by scanning a laser beam across the eye in six phases; horizontal, vertical, circular and positioning phases. Each sweep of the beam across the eye provides an image of reflectivity of the eye along a line. The laser used is a 650nm 1mW red laser attenuated to 0.15mW and the beam is manipulated by a mirror that steers the beam across the eye. The reflected light is received by a photodiode array arranged in two pairs close to the eye.

The user must wear a headset with the entire apparatus mounted and the system is limited to recording the movements of one eye only.

The system appears to work well in recording eye movements, as opposed to being used as an interface device. It works well in a clinical setting documenting eye movement abnormalities, and it demonstrates that it is possible to have a continuous scanning laser illuminating the eye at eye-safe levels.

Its primary purpose seems to be for recording eye movements rather than controlling devices. Having the user wear a headset is obviously a huge disadvantage for applications in a ubicom environment.

A different scheme for deducing gaze direction was reported by Hunter et al [17] using retinal birefringence, that is the ability of retinal fibres to change the polarity state of light. This research comes from the clinical medicine arena and demonstrates the use of lasers in

ophthalmology. Put simply, light reflected off the retina when the fovea is at the center of the scanned area has a markedly different polarization than when the fovea is not at the center.¹ Much of this description is quoted directly from their paper.

An overview of the system as shown in their paper is given in figure 2.1. The system uses a 13mW 830nm polarised laser-diode laser. The laser light is diverged by a concave lens and reflected by a nonpolarising beam splitter toward a small, rotating, concave gold mirror. The gold mirror is tilted such that it wobbles as the motor rotates. Light from the small mirror is imaged onto a large concave mirror such that, when the subject looks at the image of the light source in the large mirror, the scanning spot of light traces a small circle on the retina. The source is thus imaged into the large concave mirror by the small wobbling mirror and is subsequently imaged into the retina by the optics of the eye. The light is reflected by the eye and reimaged by the autoconjugacy of the optical system back toward the light source. Part of the light returns through the non-polarising beam splitter and enters the detection system. This is a polarising beam splitter which splits the beam into two orthogonally polarised components focuses onto two photodetectors (X and Y). The output of the X photodetector is subtracted from the Y photodetector to detect the polarisation. It was found that by scanning the retina at a frequency of 44Hz, with central fixation (the scanning beam centred on the fovea) the differential polarisation signal had a prominent component at 88Hz. With paracentral fixation, the output of the detectors was predominantly 44Hz. The differences in output frequency were more readily detected by power spectrum analysis where the power at 88Hz was over 64% during central fixation and less than 19% during paracentral fixation.

The authors also report that the movement of the eye correlated with phase shift, so it was possible to match phase variances to physical movement of the eye. This work is interesting in that it again proves that it is possible to use lasers for gaze detection and is the first system to be device-centric - i.e. it asks the question “is the user looking at the device”,

¹The fovea is the area at the back of the eye on the retina onto which the lens focuses the image the person is looking at.

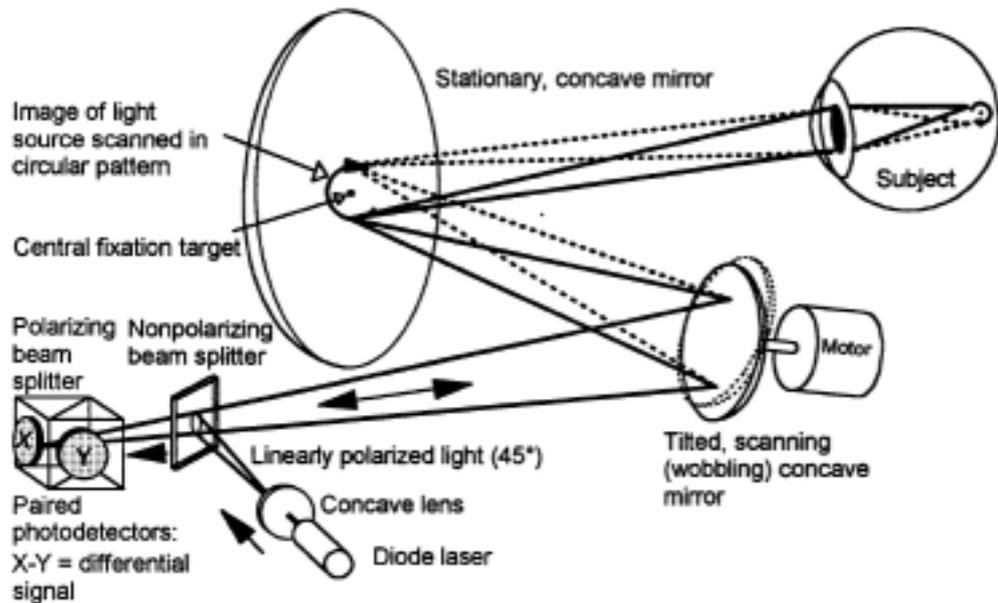


Fig. 2.1: RBS system overview

rather than “where is the user looking”. The authors say that this work could be used to operate appliances using RBS-based fixation monitors, but another beam-splitter would be needed if the apparatus including the stationary concave mirror could not be positioned in front of the user. The beam is also scanning quite a small area and the user has to be directly in front of the beam.

2.2 Related Work

Vertegaal, Shell et al at the Human Media Lab, Queen’s University, Ontario, Canada are working on a number of projects to do with eye contact and gaze detection. In [37] they designed an Eye Contact Sensor that uses a video camera and an on-axis infra-red light. This is another example of a device-centric approach, but is considerably simpler than the Hunter approach.

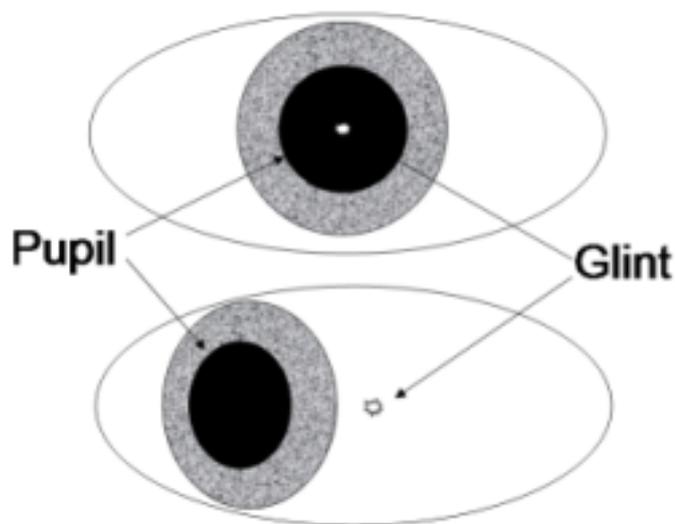


Fig. 2.2: Eye and reflected glints

The Eye Contact Sensor (ECS) works by using a video camera to locate the users pupils - at the same time it is emitting an infra-red light. The infra-red light produces a “glint” on the outside of the cornea (a reflection of the light source). When this glint is detected within the pupil area, the user must be looking at the light source, and hence the device. Even if this glint is not within the pupil area, the device can still use this information that the users gaze is “close”. Figure 2.2 shows an example of the eye (top) looking at the camera, with the glint in the pupil area, and an example eye (bottom) looking away from the camera.

The authors developed EyePliances, devices embedded with Eye Contact Sensors. The sensors monitor the environment in their field of view, searching for eye contact. When eye contact is detected, the EyePliance responds by paying attention to the user. For speech enabled EyePliances, paying attention involves listening, or activating its speech lexicon. Thus the target of speech in a single user scenario is the device that the user looks at. They built an Attentive Lava Lamp [31] which responds to the phrases “turn on/off” only when it receives eye-gaze from a user. They also built an Attentive Television [31] which pauses its

feed when nobody is watching it.

The authors have developed another system called ECSGlasses. Because the Eye Contact Sensor is pulsing infra-red illumination, the ECSGlasses can detect this. It will be possible to subtly vary the pulsing of the Eye Contact Sensor to produce a unique infra-red ID. Likewise, the ECSGlasses could also produce infra-red pulses which would enable mutual identification of device and user.

Smith, Vertegaal and Sohn also designed “ViewPointer” [32], a wearable eye contact sensor that detects infra-red tags embedded in ubiquitous devices. The user wears a small camera mounted on an earpiece that sits in front and below one eye. Because the camera is mounted so close to the eye, there is little image processing to be done to “find” the eye and simple thresholding techniques can be used to extract the dark pupil and infra-red glint. There is also no need for calibration. The disadvantage is that the user has to consciously wear a device.

The infra-red tags are very small, 1.5cm in size and height. They emit no visible light and therefore are unobtrusive. They provide similar functionality to RFID tags but with the added benefit that recognition is directional.

There are issues with the amount of data that can be transmitted vs how long a user has to gaze at the device. Larger ID numbers means longer transmit times and so the user must hold gaze for the tag to identify itself. They calculate effective transmission rates of 14 bits per second. It is also possible to transmit data rather than just an ID number. For instance a tag could transmit a URL. Using some intelligence it is possible to transmit a URL in 4 seconds. Ultimately Java applets could be downloaded and executed just by looking at an object containing an infra-red tag.

A question that arises from using these type of wearable devices is what happens in an environment when one person is wearing the device and another is not? Is the person wearing a device rewarded with a richer experience, or is the person not wearing a device penalised? This will have implications for the wider acceptance of such technology.

Chapter 3

Design

3.1 Introduction

Previous chapters have covered the concept of gaze-tracking and reviewed the state of the art. The problems of context have been discussed with regards to ubiquitous computing. This project attempts to solve one single problem: indicating to a device that a person is looking at it. It does not need to find out where the person is looking, nor indeed to track their gaze. It attempts to do this as simply, effectively and cheaply as possible. The author believes that this would assist other ubicom technologies such as voice recognition to operate with much more efficiency.

The project attempts to prove the concept that a scanning laser beam can be used to detect the presence of a human eye looking at a device. It will not use a real person in the testing but will attempt to simulate the properties of a human eye. The concept is that a laser is mounted on a ubicom device, it scans a field in front of the device and if an eye is present, reflects the light back to a sensor which is mounted close to the emitter. Only if the person is looking directly at the device will the laser be reflected.

If an “eye” is detected gazing, the system will make a simple indication such as lighting

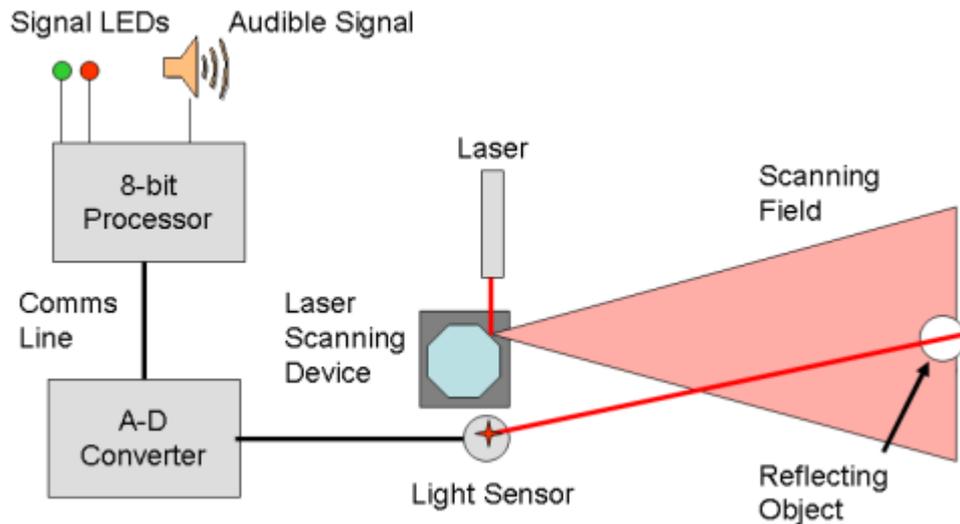


Fig. 3.1: Original design overview

an LED. In a production system, this indication could include transmitting a data packet over ethernet or even Bluetooth.

The system must not rely on predefined positioning of the target device - e.g. if it was attached to a CD player it should be possible to move the CD player without reprogramming or even notifying the system. The only requirement here is that the detection system be co-located with the device.

It is intended that this project will build a prototype “sub-system” i.e. it will be embedded in a host system. Figure 3.1 shows a high-level design diagram for the prototype, and this chapter will discuss each of the elements that will be included in the prototype.

In designing a system to work in a ubicom environment, there are some design criteria that are specific to that domain:

- The system must be low-cost. If something is to be ubiquitous and embedded in “everything” it has to be cheap to make and install in everything.

- The system must be fault-tolerant. Ubicom systems may have to run without maintenance or updates, sometimes for their entire lifetime.
- The system must be low-power. From a green perspective, more ubicom systems means more power consumption. This adds cost to the environment and to the owner. Also, ubicom systems may have to operate disconnected from mains power for periods and low-consumption design facilitates use of compact batteries.
- The system must be safe. This prototype uses visible light lasers - the production model may use infra-red lasers so extensive testing must be done to deliver a system that is verifiably eye-safe.
- The system must be easy to use and require no special training.
- The system must be secure. The prototype must be evaluated against perceived security threats. For instance, if the system is to be used in a security “call-for-help” system, how could that be compromised by an intruder?
- The system must fail safely. This will become an embedded sub-system. It must have some way of communicating to its host device that it is operating satisfactorily.

3.2 Design for Test

When developing software based systems, the developer usually has a host of debugging and simulation tools to use in building the system. With hardware based systems, there is usually a minimum of debugging facilities available. Therefore it becomes vitally important to build in testing facilities at the design stage. A modular approach should be used ensuring components can be tested individually and at any stage during the implementation process. This allows for the elimination of suspicious components during debugging. Feedback mechanisms should also be designed in to allow components show their status during operation.

Because a laser is used, it is necessary to discuss a little about how lasers work , the safety implications of using lasers for both operator and subjects, and the importance of laser safety procedures.¹

3.3 Laser technology

The word “laser” stands for “light amplification by stimulated emission of radiation”. The ruby laser was the first laser invented in 1960 by Theodore Maiman at Hughes Laboratories in Malibu, California [23].

A laser requires three things: a lasing medium, an energy source or pump, and a resonator (typically a pair of mirrors at opposite ends of the lasing medium, on which is a full mirror, the other is half-mirrored). When an energy source is applied to the lasing medium (either intense light or electrical charge) the atoms move from their base energy state to an excited energy state. After a short time, the molecules decay to their base energy state but in so doing emit a photon. Some molecules emit their photon on course parallel to the two resonator mirrors. They begin to “bounce” between the mirrors and also collide with other excited molecules triggering them to move from the excited state to base state, emitting another photon. This process causes an “avalanche” of photons bouncing between the two mirrors until eventually the light must escape - through the half-mirror - and we see a laser beam. (See figure 3.2.)

This light has some useful properties - it is monochromatic (one wavelength) and coherent. It is also either well collimated (nearly parallel rays for most lasers including gas and solid state types) or appears to originate from a point source (diode lasers) [13].

¹It should be noted that all laser safety procedures required by Dublin University were observed during this project

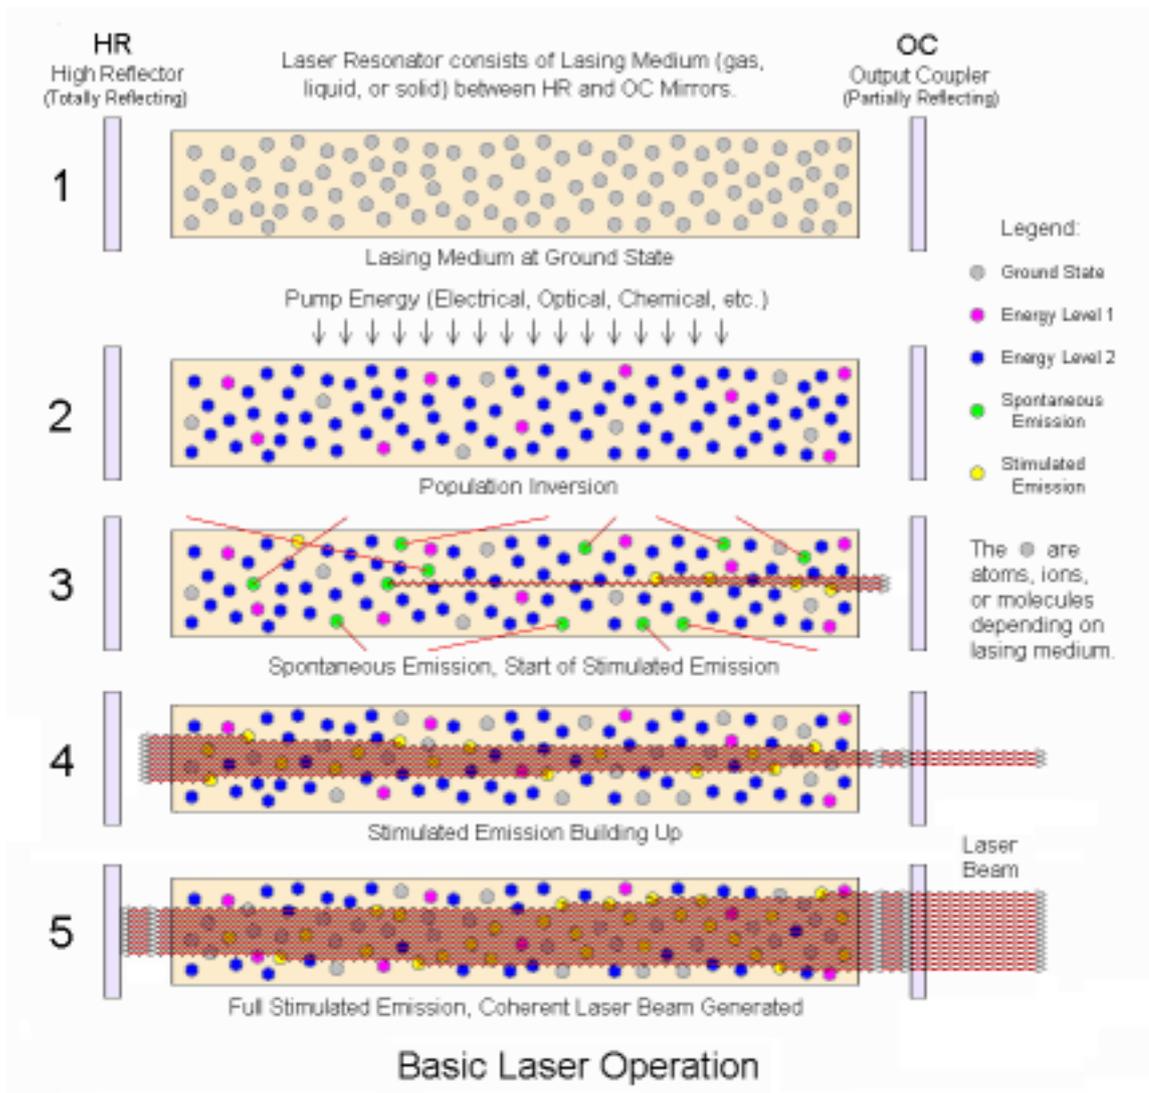


Fig. 3.2: Basic Laser Operation (Image Copyright 1994-2006, Samuel M. Goldwasser)

3.4 Laser Safety

The actual energy or power emitted by the vast majority of lasers is often quite modest. However, the spectral brightness is many orders of magnitude greater than any other known conventional non-laser light source. The intense beam can give rise to biological hazards, even for low-power lasers.

This section will cover the safety issues related to the user of lasers during the building of the prototype. These issues are also of note if the prototype is ever to be commercialised. A commercial product will have to be deemed “eye-safe” for general use in a domestic environment.

Harm from laser radiation results from the absorption of energy from the incident beam by the biological tissue which is irradiated. This energy manifests as heat. If the tissue cannot adequately dissipate the energy, damage will result. The amount of energy delivered depends on a number of factors including whether the beam is continuous wave (CW) or pulsed; its wavelength; how the energy is transmitted, scattered and absorbed by the materials through which it passes; the area of the beam at the irradiation site and energy density which is encountered there.

The eye is particularly vulnerable to biological damage from lasers because of the nature of its operation: it is built to focus visible light onto a highly sensitive part of the eye called the *fovea*. The fovea is the part of the eye onto which the lens focuses the image we are looking at. It is the only part of the eye that can see detail. Our peripheral vision is blurred. If a laser beam enters the eye via the lens, it will be focused onto the fovea. The “gain” in this case could be as high as 10^5 . Damage to this part of the eye could render a person legally blind.

Even if the lens of the eye does not focus the beam enough to do damage, viewing the beam through binoculars, telescopes or other optical focusing devices could focus the beam to dangerous levels.

The eye has a built-in protection system - the blink reflex. Some lasers can be designated “eye-safe” partly because the blink reflex engages fast enough to shut out the danger [28]. However, with infra-red or ultra-violet lasers the danger is somewhat different. Even though the lens might not be focusing the laser onto the fovea, the different wavelengths can affect different parts of the eye due to their ability to absorb or pass different wavelengths.

For instance, far infra-red and ultra-violet light is absorbed by the cornea and sclera (white of the eye) and can cause scarring. Some ultra-violet frequencies are absorbed by the lens and visible and near-infra red light is passed straight through to the retina.

Lasers, therefore, are potentially hazardous devices which require special safety considerations when being used. The UK Committee of Vice Chancellors and Principals produce a series of safety guidelines for universities [1], including one for the use of lasers which is used as the base reference in this section. These guidelines and the international standard IEC 60825-1 (as implemented in Ireland in IS/EN 60825-1 [2]) form the basis for laser safety procedures in Dublin University as documented in the Laser Safety Training manual [7].

Lasers are classified according to their power, design and accessible emission limits. The classes are described in

The two types of lasers that this project uses are a Class 3A laser-diode laser and a Class 3B** Helium-Neon laser.

A Helium-Neon laser is one of the most common lasers and is a general workhorse in laboratory environments [6]. It consists of a glass tube filled with a mixture of helium and neon gas. The electrodes are mounted inside the tube which is sealed. The mirrors are also bonded to the tube and are not adjustable. HeNe lasers have a long life time and usually produce light at a wavelength of exactly 632.8nm. This is important as most laboratory experiments are working with specific wavelengths of light and the HeNe laser produces a predictable wavelength.

A laser-diode laser is a solid state laser which uses a p-n junction as the lasing element [6]. It is extremely compact, power efficient and cheap. The most visible application of this type

Class	Type	Description
Class 1	Safe	Either inherently safe or safe by design
Class 2	Low Power Visible CW and pulsed	Protection afforded by natural aversion responses including blink reflex. Hazard controlled by simple procedures
Class 3A	Low/Medium Power	As with C2 but viewing with optical aids may be hazardous
Class 3B*	Low/Medium Power	As with 3A but slight hazard from viewing beam directly
Class 3B**	Medium Power	Hazard from direct beam viewing
Class 4	High Power	Not only a hazard from direct viewing but also from specular reflections and possible diffuse reflections. Extreme caution

Table 3.1: Laser Classes

of laser is the laser pointer, but it is also commonly used as alignment guides for industrial cutting equipment. In this case, calibration of the wavelength is not important as long as it is visible. This causes problems when trying to filter the laser light to reduce ambient light detection as low-cost lasers can have variations on the specified emitted wavelength of $\pm 10\text{nm}$.

The safety requirements for a Class 3A laser are [7]:

Beam Path: Terminate beam path at useful length

Training: Required for all operators

The safety requirements for a Class 3B** laser are [7]:

Remote Interlock: Connect to room or door circuits

Key Control: Remove key when not in use

Beam Stop: Prevents inadvertent exposure

Emission Indicator: Indicates laser is energised

Warning Signs: Posted on door etc

Beam Path: Terminate beam at end of useful length

Specular Reflections: Prevent unintentional reflections

Eye Protection: Required

Training: Required for all operators

3.5 Red-eye and retroreflectors

The red-eye effect has been mentioned a number of times and it is relevant to this project to discuss how it works and why its important to the project. Red-eye is sometime seen in photographs where the subjects have an eerie red glow in their eyes [30]. This is caused by the light from a camera flash reflecting off the retina at the back of the eye and being picked up by the camera. The reflected light is red because the retina is covered with blood

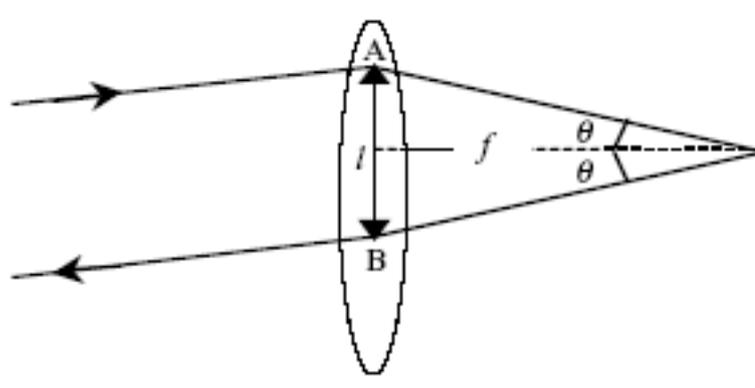


Fig. 3.3: Retroreflector Principle

vessels which tint the colour. The light is seen quite strongly because the eye (in this case) is acting as a retroreflector. A retroreflector is a device that reflects a light source back to its origin - either along the exact same path or on a parallel path close to the original path. A simple retroreflector can be made using a lens with a mirror mounted exactly at or close to its focal point. In 3.3 the light entering the lens at A is focused onto the focal plane at angle θ . The light then reflects off the focal plane, again at angle θ and is refracted at B back to a parallel course [26]. The most common retroreflector we come across are the “cat’s eyes” we see mounted in the center and sides of roadways.

As already mentioned, a retroreflector is simply a lens with a mirror mounted exactly at or close to its focal point . A “cat’s eye” uses a simple glass sphere as a lens with silvering mounted on the rear half of the sphere to act as a mirror. The lens in our eyes focuses anything we look at onto the retina and so if the flash on the camera is too near the camera lens, our eye lens focuses the flash onto the retina, which reflects it back through the eye lens and back toward the camera. Only the camera “sees” the red light as it is traveling back on a closely parallel path.

The retina in a human eye is not designed to be reflective so much light will be lost. Some

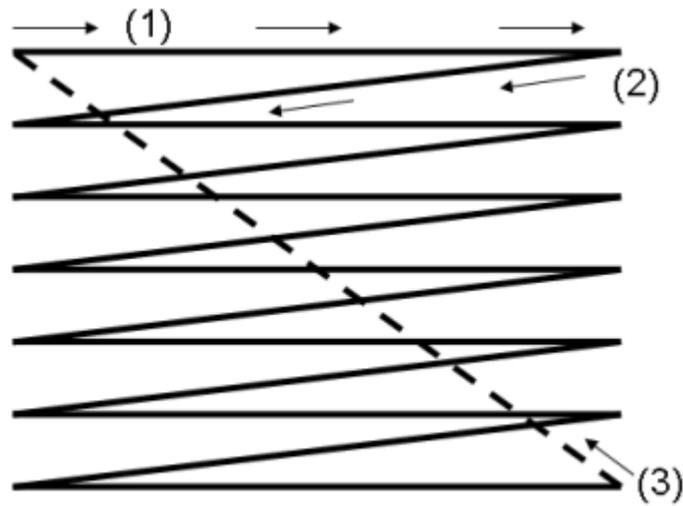


Fig. 3.4: Raster Scan

animals, particularly cats have a reflective layer called the *tapetus lucidum* which is designed to reflect incoming light back through the retina, effectively firing the light sensors in the animal's eyes twice increasing the it's low-light vision.

3.6 Scanning

A laser emits a fixed beam of coherent light in one direction. If it is to be used to detect a human eye, it has to be manipulated so that it “scans” a zone in front of the ubiquitous device. The original idea was that the laser would have to be scanned in a similar way to a raster scan.

Looking at 3.4, in a raster scan, the laser would scan from left to right (1), then move down to the next scan position (2) and so on to the end of the scan zone and then return to the start position (3).

Typical laser scanning devices use a closed-loop galvanometer and mirrors to manipulate

the beam. Its easier to move the beam rather than the laser.

3.7 Light detector

Obviously if the laser emits a beam, a detector is needed to pick up the reflected beam. A silicon detector provides a means of transforming light energy to an electrical current, using the photovoltaic effect (similar to a solar cell). Because the laser is coherent, the beam that leaves the emitter and is reflected off the retina, will return somewhat intact and fall onto the detector, thereby triggering an increase in current from the detector.

3.8 Analogue to Digital converter

The light detector is an analogue device so an A-D converter is required to turn its readings into discrete digital values that can be queried and analysed by the processor. There is only one signal to process so a simple one-channel A-D converter is required.

3.9 Processor

The system only has to detect if a person is gazing at the device and make a simple indication, so the amount of processing required is minimal. In keeping with the aim of making this system cost effective, an 8-bit processor such as an 8051 will suffice. These devices now typically cost a few cents. They are also available in many different packages, eg with on-board ethernet controllers, so the initial development could be modified to take advantage of any 8051 package.

Chapter 4

Implementation

4.1 Introduction

The prototype was built with “design-for-test” in mind. The prototype is primarily a hardware device and trying to debug hardware issues can be immensely time-consuming. Therefore components were selected to enable rapid fault-finding and diagnosis ability.

During the build phase, it became clear that the laser originally selected would not be able to perform all tests. The original laser was a 5mW laser-diode laser. A second laser was acquired, a 10mW Helium-Neon (HeNe) laser and the evaluation tests split into two sections.

One interesting fact that came out of testing was that the original premise that the eye was a natural retroreflector because of the spherical nature of the eye-ball itself was proven false. The eye acts as a retroreflector because the lens of the eye focuses light on the retina, which reflects bright light back through the lens to the source.

Figure 4.1 shows an overview of the prototype with the two processor boards at the bottom right of the picture, and the two reflector targets at the left. The long tube in the centre is the Helium-Neon (HeNe) laser, pointing at the silver, octagonal rotating polygon. Just below the polygon and dwarfed by the HeNe laser is the laser-diode laser. To the right

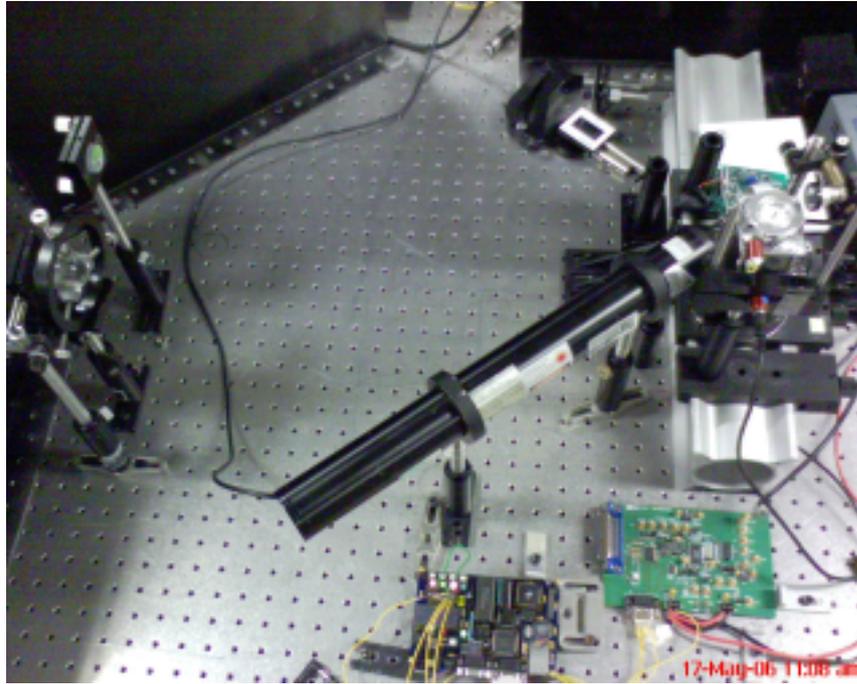


Fig. 4.1: Overview of the prototype

of the polygon is the light sensor. The equipment is set up on an optical bench and black shields have been erected around the prototype to ensure the laser light is contained.

4.2 Hardware

4.2.1 Processor - RITA51J Evaluation Board

The processor selected for this project was the 8-bit 8051. Any 8051-based package would have sufficed, but in designing-for-test, an evaluation board package was selected. Evaluation board packages allow developers build test systems quickly as the board includes both hardware and software environments for the the processor chip to operate in. For instance, most have built in RS-232 ports, onboard RAM and ROM, and some form or Integrated

Development Environment (IDE) to allow the developer program in a high level language like C and then manage the download the chip.

The RITA51J from Rigel Corp [5] (as seen in figure 4.2), is an evaluation board intended for Internet and intranet applications. As such it has most of the properties of other commonly available evaluation boards, but it includes a CS8900A ethernet LAN controller. Although this project does not use the ethernet adapter, there were a number of features that made this board more useful than similar packages:

1. It uses a physical switch to switch between monitor mode and operation mode, so that the same serial port can be used for both downloading and I/O
2. All system signals are available via headers
3. It has 4 leds and 4 switches installed on the board that can be patched to appropriate header pins. This means the LEDs can be used for debugging and reporting program status and ultimately to represent sample actions by the program.
4. Rigel Corp provide a library of useful serial port routines which assisted greatly in debugging
5. There is no limit on the size of the program produced by the READS51 compiler. Some companies give an IDE away free with the board, but limit the compile size to 2K. The “full” version must be purchased to enable builds greater than 2K.
6. It has 128K RAM and 128K on-board FLASH ROM

The IDE that Rigel Corp provides includes a C compiler that is based on Small-C [16]. Small-C implements a subset of the Kernighan and Richie C language. It was written by Ron Cain and published in the May 1980 issue of Dr. Dobbs’ Journal. Later, James E. Hendrix improved and extended the original Small-C compiler. Originally written to produce 8080

assembly language it has been ported to several processors and microcontrollers. It has some limitations compared to K&R C or ANSI C including:

- Structures and unions are not implemented
- Only one-dimensional arrays are allowed
- Only one level of indirection is allowed
- Only integer and char types are allowed.

4.2.2 Analog to Digital Converter - AD7718

The AD7718 analog to digital converter [11] (as seen in figure 4.3) is a low-voltage, low-power, sigma-delta analog to digital converter designed for low frequency measurement applications. It can be configured with 8/10 pseudo-differential input channels or 4/5 fully-differential channels. Input signals can range from 20mV to 2.56V and it operates off a 3V or 5V power supply. The AD7718 supports three-wire SPI communications protocol.

The AD7718 also comes in an evaluation board package, the EVAL-AD7718-EB [10]. This package mounts the AD7718 chip on a circuit board and breaks all analog input pins out to micro-BNC plugs. It also maps the I/O pins to both a 9-pin serial port and a Centronics-type parallel port. This latter port is used by the software provided to program and analyse the AD7718 via a standard PC. The board also has two AD780 precision 2.5V references, a 32.768kHz crystal and digital buffers to buffer signals to and from the PC.

The AD7718 employs a sigma-delta conversion technique to realise up to 24 bits of no missing codes performance.

In keeping with the “design-for-test” principle, this evaluation board package allowed the light sensor to be tested without the need for the processor to be programmed, and for the best mode to be selected. It also allowed for confirming data transmitted from the AD7718 to the 8051 during debugging.

The AD7718 was tested in a number of modes, but the optimal modes for this project were deemed to be:

- 2.56v input
- Unipolar mode (0 to $1.024V_{ref}/\text{Gain}$ represented by 0000Hex to FFFFFFFHex)
- Default hex filter value 45hex
- Pseudo-differential channel
- Using Ain5 and Aincom (analogue input 5 and common reference)
- Chop mode enabled

The reason 2.56v mode was chosen was because of the limitations of Small-C. Small-C only supports CHAR or INT data types. The INT data type is implemented in Small-C in 16bits. The program on the 8051 had to be able to calculate average ambient light values to create a “floor” value. As can be seen from table 4.1 below the maximum value output by the sensor was 446mw. However, adding multiple AD output values at this range would have overflowed the 16bit INT. It was simpler to reduce the range of the AD7718 by setting the input value to 2.56v.

This allowed the AD7718 to operate with an ADC update rate of 19.79Hz producing updates every 50.34ms. The output coding of the AD7718 in unipolar mode is:

$$Code = \left(AIN \times GAIN \times 2^{24} \right) \div (1.024 \times V_{REF})$$

4.2.3 Rotating Polygon Laser Scanner

The initial design of the prototype called for the laser to operator a raster-scan across the front of the ubicom device. This would have meant trying to maneuver the emitted beam

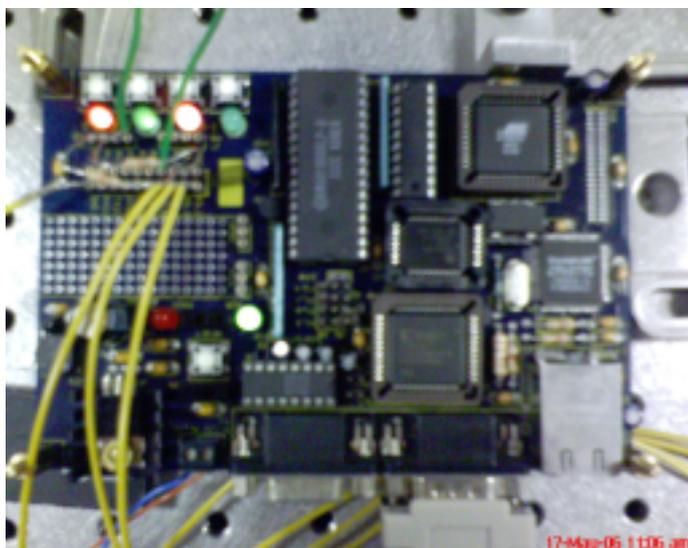


Fig. 4.2: Rita51J evaluation board

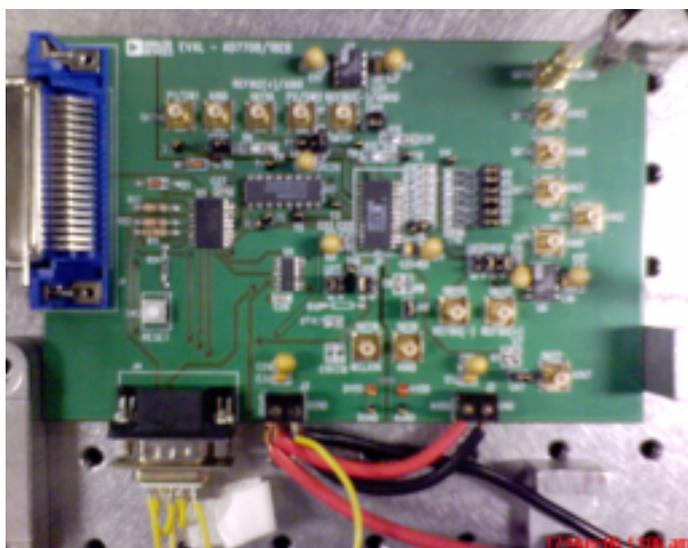


Fig. 4.3: AD7718 evaluation boards

in two dimensions. However, lasers used in alignment tasks sometimes have a lens attached which turns the “dot” into a line. This is called a “fan-angle” lens. If a laser of this type is used, then the line can be reflected off a simple rotating polygon scanner and so generate the “field” of light in front of the ubicom device. Many of these scanners are built to high tolerances but in this case tolerance is not important, so a relatively low cost package can be used.

Lincoln Laser produce rotating polygon components and also supply three different standard packages. The SOS-BB10 was used in this prototype [21] (as seen in figure 4.4). This package consists of the rotating polygon itself, a motor and controller. This device can run from 1,000 to 10,000 RPM and has a reflectivity index of 88% at 633nm. (12% of the laser light is lost in the reflection). With 8 facets to the polygon, that represents 133 scans per second at 1,000 RPM (133Hz) giving a scan time of 7.5ms. At 10,000 rpm the scanner produces 1,333 scans per second (1.3kHz) with a scan time of 0.75ms. This package does not report the position of the polygon at any time (so determining where exactly the beam is is not possible) but this was not necessary for the project.

4.2.4 Light Detector

The light detector chosen was a silicon photovoltaic detector from Edmund Optics [3]. This detector responds to visible light, has a large surface area (100mm²) and a fast rise time of 43nsec at 632nm (time taken for the sensors output to go from 10% to 90% of its final value). It has a bnc connection.

4.2.5 Lasers

The initial laser chosen was a small, cost-effective laser-diode laser from Global Laser, the LaserLyte V3VLL [35]. This is a 5mW visible red-light laser that uses a built-in lens (fan-angle) to produce a line rather than a dot. The light emitted is 635nm \pm 10nm. This

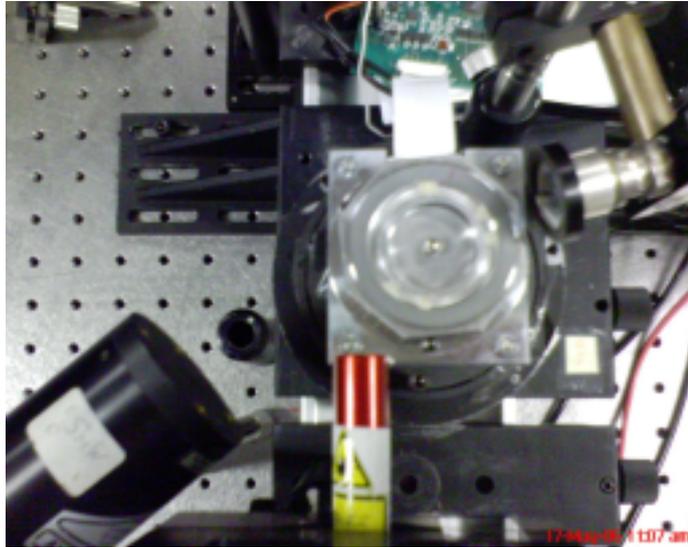


Fig. 4.4: Rotating Polygon Scanner

laser is primarily used for industrial alignment operations such as setting up a fabric cutter. Consequently it is more important for this laser to produce bright visible light rather than be calibrated to a specific frequency. Most lasers in laboratory use are calibrated to one frequency and usually to a common specification (for example the HeNe laser discussed below). Both the design of the laser-diode laser and this lack of a requirement for wavelength calibration means these lasers are mass-produced, compact and inexpensive.

A more powerful laser was used in the latter half of the project. This was a 10 mW Helium-Neon laser Model 1125 from JDS Uniphase [9] producing a dot at 632.8nm.

4.3 Software

4.3.1 Supplied Software

The software used in the project includes both software supplied by the manufacturers of the equipment and software written for the 8051 board. The software supplied with the Rita51J

board is an integrated development environment (IDE) which includes a compiler based on SmallC. This allows for developing applications in a high-level language and then manages the downloading and communication with the monitor on the evaluation board. The IDE includes a terminal emulator which allows for direct interaction with the on-board monitor for controlling download, debugging and erasing the flash memory. The IDE is provided with the board and unlike some other packages does not have an evaluation mode which limits program compile size. The software supplied with the AD7718 allows for programming the AD7718 itself and reading the data returned from the AD converter. This was important for testing the AD7718 and for performing evaluation tests on the data generated by the light sensor without having to have the 8051 on-line.

4.3.2 Written Software

The AD7718 communicates with an outside processor using the Serial Peripheral Interface (SPI) protocol [20]. SPI is a serial bus standard designed by Motorola and supported by various manufacturers. It is a synchronous, serial data link that operates in full-duplex mode. Devices communicate using a master/slave relationship with the master initiating the data transfer. There is no formal packet protocol, nor is there a formal acknowledgment scheme - it is up to the communicating parties to understand the transactions. There is also no flow-control. Consequently, it can be tremendously efficient for on-board communications between components. Some 8051 packages such as the Keil evaluation board, include an on-board SPI processor however the Rita51J does not. Therefore code to manage the SPI communication to the AD7718 had to be written.¹ The application part of the code is very simple; it samples the incoming data from the AD7718 and turns LEDs on or off depending on the data returned. There are three LEDs used as indications in the prototype:

- LED 1 is lit to indicate that the software has completed its initialisation and has

¹The author is indebted to John D. Maniraj for providing sample 8051 code

calibrated itself to ambient light levels.

- LED 2 is lit when the software determines that it is being looked at (eg light is reflected from the “retina” onto the sensor). It remains lit while it is being gazed at.
- LED 3 is used as a simple application demonstration as explained below.

LED 3 is used to demonstrate potential applications for the prototype. The idea is that if the reflection is seen for more than a second (ie the user holds his gaze at the device) this LED is lit and holds its state. This represents a user using his gaze to directly affect the device eg power it on. If the reflection is seen again for more than a second, the state flips eg the LED turns off. If the gaze is held continuously the LED alternates on and off, indicating that the system could be cycling through menu options, and breaking gaze selects the option.

4.4 Building the prototype

The prototype was built in one of the Physics Department laser laboratories in the SNIAM building. This allowed for a controlled, safe environment to experiment with lasers. The lasers, sensors and boards were all fixed to an optical bench. Two target reflectors were used in the testing - these can be seen in figure 4.5. The first was a mirror which was slightly concave (on the right of the picture). As the scanning column passed the reflector the light was reflected into the light sensor. This mirror was used to test the if the field of light projected by the scanner and laser with the fan-angle could be successfully reflected and detected.

The second target was a quartz sphere (seen on the left of figure 4.5), with a flat mirror behind it. Initially, the quartz sphere was acquired in order to approximate the eye, but subsequently was used at a simple lens. The mirror was set behind the sphere at the focal point of the sphere to reflect the beam back through the lens. Ideally this should have been a spherical mirror but this was sufficient for testing. This provided the retroreflector to test if

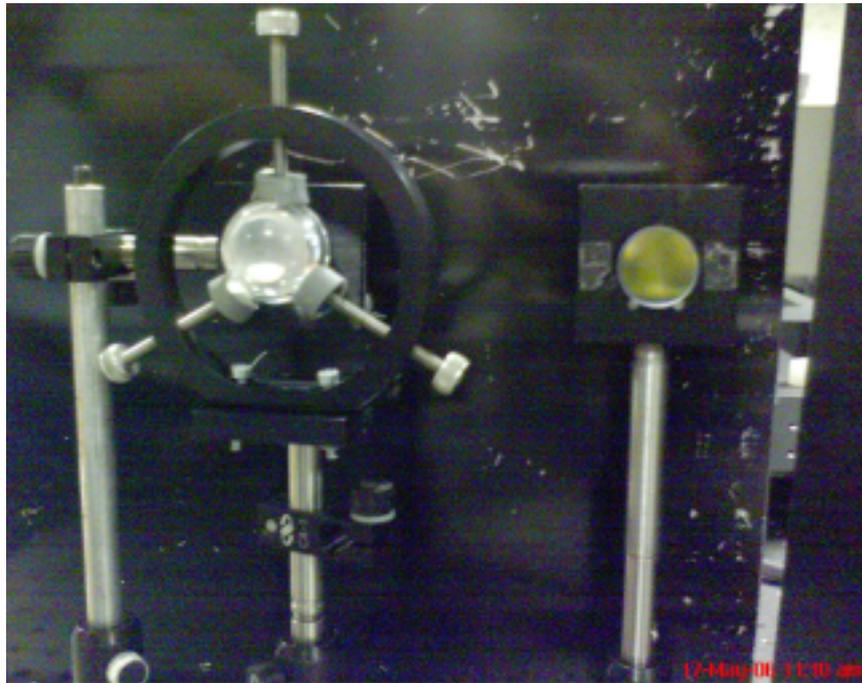


Fig. 4.5: Two reflector targets

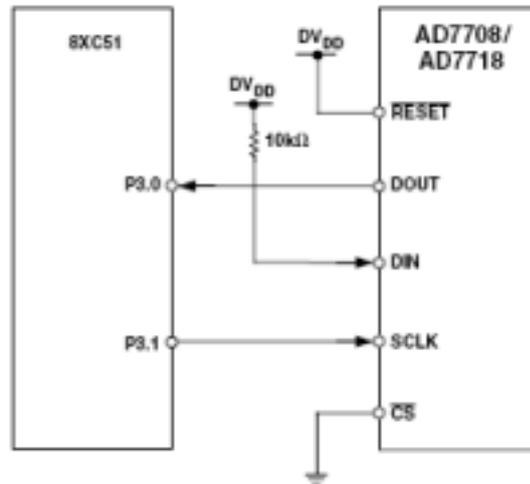


Fig. 4.6: Original SPI interconnect

light could be reflected by and “eye” and detected. Both lasers were mounted to operate with the scanner and the light sensor was mounted close to the scanner (considered the originating point for the scanning beam). Each mirror was covered while the other was being tested.

4.4.1 Interconnecting the boards

The two boards were connected using SPI via the 9-pin port on the AD7718 and the P1 header on the 8051. The grounds of both boards had to be connected together to provide a common ground for the comms signals. In a production system, both chips would be on the same circuit board with a common power supply.

The documentation for the AD7718 gives a diagram for connecting to an 8051 as shown in 4.6. The connections as actually implemented are show in 4.7.

The reason for this change is that the original specification [11] called for the interconnection to be made via pins P3.0 and P3.1 on the 8051 and the 8051 UART to be set to Mode 0 (single line I/O). These pins are mapped to the TX and RX channels on the UART.

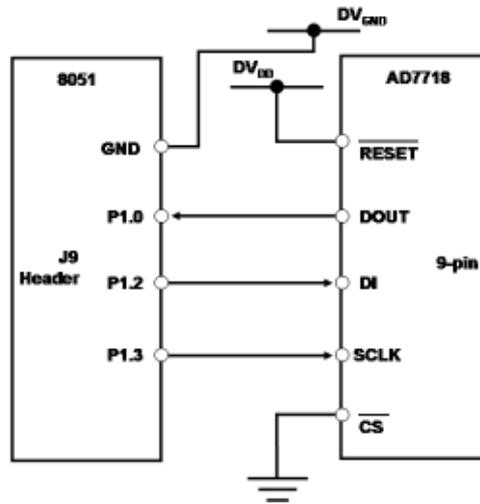


Fig. 4.7: Final connections

However, the project needed to use the serial port to communicate with the IDE during development and so pins P1.0 and P1.1 were used. However, the P1 port on the 8051 implements quasi-bidirectional output. This means the port can be used for input or output without the need for reconfiguring. It can do this because when the port outputs a logic high, it is weakly driven, allowing an external device to pull the pin low. Simply put, an external device can override the signal being generated by the 8051. In the original diagram the DOUT and DIN pins of the AD7718 are connected together and pinned to a 10kohm pull-up resistor. This provides a two-wire connection and is considered the most elegant and simple. But because the connection was now with P1.0 and P1.1 this meant that the DOUT pin was holding both P1.0 and P1.1 high. The output was being latched by the 8051 but being held high by DOUT. The solution was to move to a three-wire solution. The grounds of both boards have to be interconnected to provide a return ground for the TX, RX and clock lines.

An overview of the software implementation of the SPI protocol as implemented in this project is given in 4.8.

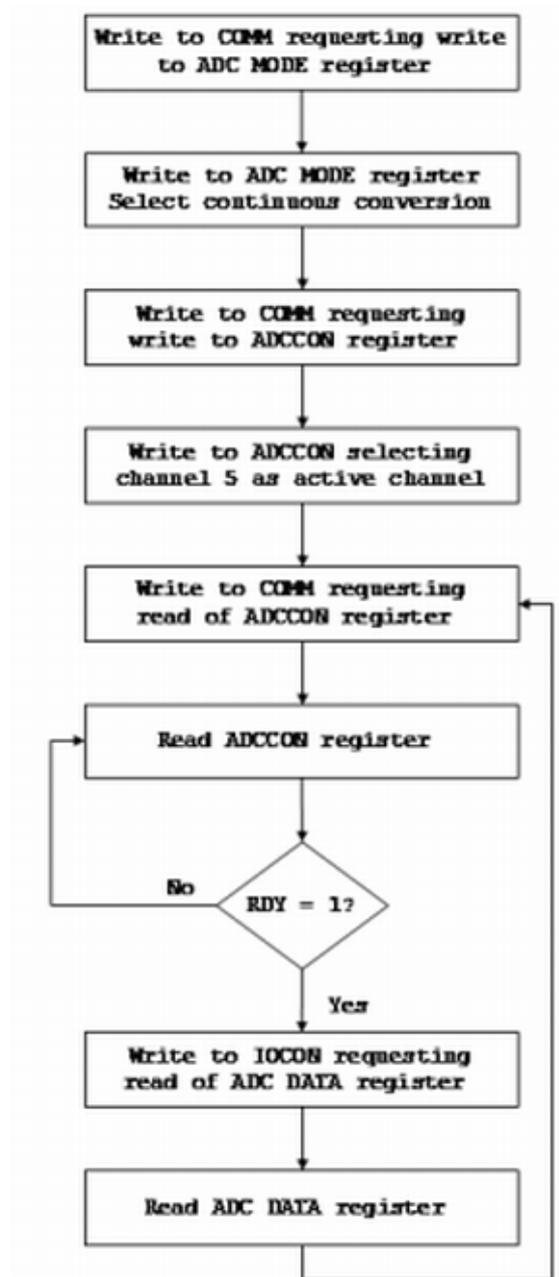


Fig. 4.8: SPI Data Register Read cycle

Source	Average Value	Voltage	V Diff	% Diff
Ambient Light	1C0331	280mV	0	
5mW Direct	---	357mV	77	27.5%
5mW Direct Reflection	217E2A	335mV	55	19.6%
5mW Scan Reflection	1C70B8	284mV	4	1.4%
10mV Direct	---	446mV	166	59.2%
10 mW Direct Reflection	2960FF	414mV	134	47.8%
10 mW Scan Reflection	1C84C2	285mV	5	1.8%

Table 4.1: AD values returned

4.4.2 AD Values in Detail

The values returned by the AD7781 are outlined in 4.1. Firstly the value for ambient light was recorded, then each laser was shone directly into the light sensor to get a “maximum” value. Then the values were taken from the laser reflecting off the scanner mirror while static, up to the target and back to the light sensor. Then finally with the laser reflecting of the scanner while rotating, up to the target and back to the light sensor.

The values in “Average Value” are the 24-bit average values from a continuous run of 200 samples of each type. The “Voltage” column shows the average voltage reported from those same samples.

Taking the ambient light as the “floor” value, the “V Diff” column is the difference in voltage seen at the light sensor. The “% Diff” column is the equivalent percentage difference. It can be seen from the voltage values that the 5mW laser loses 39% of its power in a direct reflection, and 95% of its power in a scanned reflection. Likewise for the 10mW laser, it loses 21% of its power in a direct reflection and 97% of its power during a scanned reflection. So the level of returned energy is very low - only 4%/5% returning to the source to be

detected. Furthermore, this experiment showed that the quartz sphere and mirror was quite inefficient in returning the light to its source. The prototype was built and tested in controlled laboratory conditions, but this small percentage would not be dependable in the real world. Further work would attempt to use filters to reduce the ambient light and only filter the returned laser light, and to experiment with more powerful lasers.

4.4.3 Timings

In trying to find out where the loss of energy is happening, it is possible that the scanner is scanning too fast and that the light sensor does not have time to complete its rise time of 43nsec (time taken for the sensor to go from 10% to 90% of its maximum output. In Figure 4.9 it can be seen that at minimum rotational rate of 1,000rpm, the laser will take 7.5ms to complete one scan (assuming a scan length of approximately 2 metres). The surface area of the light sensor is 100mm (10mm x 10mm) so the reflected laser will take 37.5nsec to traverse the surface of the sensor. This is under the rise time of the sensor of 43nsec. So the sensor will not be able to output its maximum value. The scanner was actually operating faster than the minimum of 1,000 rpm but it is not possible to determine exactly the rate as it is set using an analog screw adjustment. So for instance at 10,000rpm the time taken for the laser to traverse the detector is only 3.75nsec which would not give the sensor time to react fully.

4.4.4 Loss of Power in the Returned Signal

The figures in Table 4.1 show a significant loss of signal returned to the light sensor when the scanner is operating - 95% of the signal is lost. Given that the signal is quite strong during direct reflection, the scanning process must be causing the loss.

Nyquist states that in order to sample a given frequency without aliasing, the sampling rate must be at least twice the sampled frequency [27]. If at the minimum rate of 1,000RPM

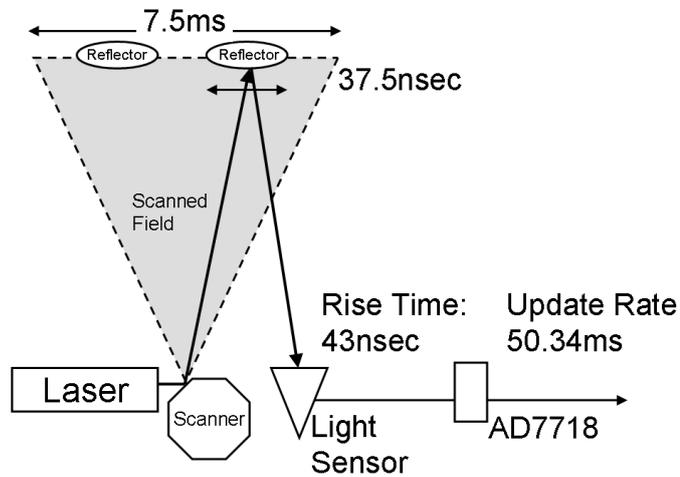


Fig. 4.9: Scanner timings

this base frequency is 133Hz therefore the sampling be at least 266Hz. The AD7718 scans the input frequencies at 32.768mHz so this is not the problem.

However the duty cycle of the light sensor is 37.5nsec which is so small in relation to the scanning rate. The AD7718 is effectively producing an average of the signal during the 7.5ms of each scan. The light sensor requires 43nsec to reach 90% of maximum output so at the slowest scan the detector is only reaching 79% of its output (assuming the output is linear).

The AD7718 configured as detailed above issues updates from the digital side at 19.79Hz or every 50.34ms. This means the scanner has scanned 6.72 times during the digital duty cycle. It is obvious then that the light sensor is spending most of its time exposed to ambient light levels and a very small time exposed to the laser as it passes during a scan. The AD7718 is reporting the average light sensed during its duty cycle.

Chapter 5

Evaluation

5.1 Design

5.1.1 Cost

The original design of the prototype was intended to be simple, to use inexpensive components and to be used to prove a concept. Both the 8051 eval board and the AD7718 AD converter worked successfully in their evaluation-mode packages, but there is no reason they could not be reduced down to a double component board. Prices for 8051 packages on the Radionics website [4] are as low as 1.25 euros for quantity purchase and for the AD7718 are as low as 10 euros.

The original Laserlyte laser was not strong enough for the test and a more powerful HeNe laser was acquired from the Physics Department to complete testing. However, the Laserlyte laser is available in a more powerful package and its benefit is its relatively low cost given it is used in industrial applications and does not require calibrating.

5.1.2 Power

Both the 8051 and the AD7718 have “low-power” modes to conserve power. Both chips could switch to low-power mode during idle. This was not implemented in the original design. The 8051 eval board required a 9V power supply, and the AD7718 a 5V supply. The laser scanner required the most power during the prototype testing (24v), but this is the only electro-mechanical device in the design. The 24v was only required when the scanner was getting up to speed, its power requirements dropped to 9V during normal operation. It is envisaged that designing a commercialisable product would require some form of non-mechanical, low power scanner. Finally, the Laserlyte scanner required a 9V power supply.

5.2 Implementation

The Rita51J 8051-based evaluation board was an excellent package to base the prototype around. The compiler had no limitations on program size, and the board had features which helped in development and testing such as on-board patchable leds, writable flash memory and switchable run/debug modes. Ultimately it has a built-in ethernet controller for extending the project onto networks

The Rita51J board operates in one of three memory architectures and one of two operating modes.

1. RIM (Rigel Industrial Mode)
2. RRM (RAM Resident Monitor)
3. RROS (ROM Resident Operating System)

Each of these architecture modes operates in either PGM (program) mode or RUN mode. The architecture mode is selected by jumpers on the J5 header, and the mode via a slide switch and reset button.

RIM architecture allows for loading code into the onboard FLASH. RRM architecture does not use ROM and RROS is used to support the assembly language debugging features of the monitor.

Switching between RROS mode and RIM mode is as easy as popping jumpers.

The AD7718 evaluation board was an excellent package to use as the AD converter. Although a simple single-channel converter was required, the eval package allowed for testing the light sensor and evaluating the best mode to set the converter to without having to develop the full application. Its eval software also allowed for verifying data returned to the 8051 board as being correct. The AD7718 eval board uses miniature BNC sensors to connect to the two analog input pins and this was a source of frustration as the light sensor used a single BNC connection.

One of the benefits of using the AD7718 eval package was being able to test and evaluate how the AD converter works and how it communicates using SPI. This effectively “ring-fenced” the AD testing from the main board. For instance, if an 8051 package with an on-board AD converter was used, not as much would have been learnt about how the AD works and interacts with a host.

The laser selected was, on reflection, not powerful enough. However, it did prove that the combination of a scanner and the laser with a fan angle did produce an acceptable “scanning field” which could reflect enough light to be detected.

The laser scanner worked adequately once a suitable power supply was found. The controller board as an enable pin which needs to be pulled low, but this can be crossed to the ground on the board to have a permanent “on” switch. It would be possible to connect this pin to the 8051 and have it control the scanner powering up and down. The scanner can be set to run from 1,000 RPM to 10,000 RPM using a screw-adjuster on the controller board.

As discussed in the previous chapter, the scanner scan rate was too fast. The laser needs to fall on the light sensor for a longer time than 37nsec. The AD7718 produces an average value during the period of the scan so work has to be done to calculate the optimum scan

time vs exposure time to produce the best average values.

The light sensor was a standard silicon detector from Edmund Optics. At the time of design the rise time of this detector was deemed adequate and the detector was chosen for its large surface area. This surface area while convenient, was not necessary. A sensor with a faster rise time would be needed to optimise the response of the returned light.

The 8051 was programmed to take five readings from the AD7718 and arrive at an average ambient light value. The difference in the ambient light vs the light reflected during operation of the scanner was very small - less than 0.2%. This meant that changes in the ambient light (including wearing a bright shirt and moving closer to the sensor) would trigger a “gaze detect” response. Using an appropriate filter would reduce the effect of ambient and random light. However, most filters for laser applications are designed for specific laser frequencies. For instance it is possible to get wide and narrow bandpass filters for the Helium-Neon laser frequency of 632.8nm. Using a filter could reduce the incoming frequencies down to a 10nm band around the base frequency of 632.8nm (628.8nm to 637.8nm).

Because the Laserlyte laser is used in industrial alignment applications, its specific bandwidth does not need to be calibrated. Using a spectroscope, the Laserlyte was found to be emitting a wavelength of 636.8nm. This would be just in the window of the HeNe narrow bandpass filter, but as the filter has a tolerance of $\pm 2\text{nm}$ it might not have worked. However, a filter would be necessary to eliminate the ambient light and to stabilise the prototype, so this would have a bearing on the choice of laser ultimately chosen to drive the prototype.

The quartz sphere and mirror also seemed to be very inefficient at reflecting the laser light. The 5mW laser was unable to return a signal using the scanner and the 10mW was just about able to register a signal but by using a normal spot beam and not a line beam. A true spherical retroreflector should have a spherical mirror behind the lens but this was untested at the time of completing the project.

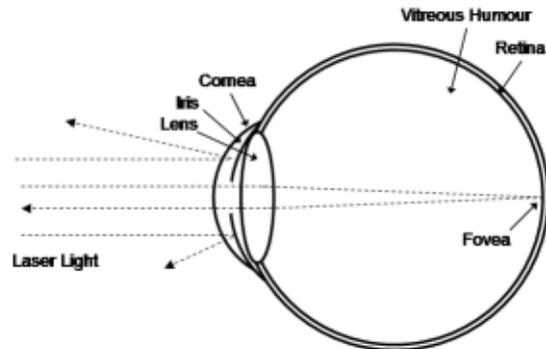


Fig. 5.1: Simplified diagram of the eye

5.3 Original Objectives

The original objectives of the project are now discussed:

Does a spherical retroreflector simulate a human eye successfully?

The author assumed because a “cat’s eye” road reflector is a glass sphere backed with silvering that that is how a real cat’s eye works and how the inventor of the road reflector got the idea. Perhaps the spherical shape of the eye, filled with the liquid *vitreous humour* acts similar to a glass sphere? Human eyes are similar so they should reflect in the same manner?

A human eye and a cat’s eye will reflect light but the cat’s eye is much more efficient. The cat has a reflective layer called the *tapedus lucidum* behind its retina, which reflects the incoming light back through the light sensors (rods and cones) so they effectively “fire” twice, increasing the cat’s night vision. The real reason a cat’s eye and a human’s eye reflect light is because the eye acts as a lenticular retroreflector (see figure 5.1).

A lenticular retroreflector only requires a convex lens and a mirror at the focal point. The road reflector uses the glass bead as a crude lens to focus the incoming light on the silver backing, not as a replica of the spherical eyeball.

Can the reflected laser energy from the retroreflector be detected by a photoreceptor, separately from the normal energy reflected from a surface?

Enough light was returned to the light sensor from both the test mirror and the quartz sphere to effect detection.

Using the 5mW laser and the scanner, enough light was returned using the mirror reflector for the light sensor to emit a detectable voltage difference, but not enough from the sphere.

Using the 10mW laser and the scanner, enough light was returned using the sphere, but this laser was not tested using a fan-angle to produce a scanning line.

Can the reflection be detected continuously enough to realise it is a gaze?

Even though the laser was scanning across the surface of the reflection, and the reflections were effectively “pulsed” onto the photoreceptor, the AD converter effectively reported continuous levels to the processor so it was possible for the processor to recognise a gaze.

What is the minimum laser power required for this simulation to work?

The author was not able to calculate the necessary minimum power, but a 5mW laser is not powerful enough but experiments pointed to possible a 15mW laser or higher would be necessary. Even if a more powerful laser is used, it is possible to reduce the power using neutral-density filters to calculate minimum power necessary to make the prototype eye-safe.

What type of photoreceptor works best?

The original photovoltaic light sensor did not have a fast enough rise time to respond to the scanned laser beam and emit the maximum voltage.

What is the best method to cause the laser to scan a raster-like pattern or a diffused “fan” pattern?

The 5mW laser with the built in fan angle produced a column which was scanned using the rotating polygon. This was easy to configure with off-the-shelf parts. A system based on closed-loop galvanometers, as an alternative, was not tested.

Calculate optimum “scan” rates

The rotating polygon scanner can operate from 1,000 RPM to 10,000 RPM. As the scan rate increased the time taken for the reflected laser beam to traverse the scanner reduced. Even at 1,000 RPM the traversal time was less than required for the sensors rise time. The optimum scan rate is therefore less than 1,000RPM.

Determine if energy required can be brought within eye-safe limits for a visible light laser.

This was not determined during the project due to time considerations.

Determine if this limitation can be implemented using an infra-red laser (infra-red lasers do not invoke the human blink-response)

This was not determined during the project due to time considerations.

5.4 Security Considerations

The current prototype could have numerous applications or client devices which respond to the gaze detected indication. Whereas the threats identified do not necessarily compromise the detector itself, the detector system could be used to compromise these applications or client systems. For instance, a possible application of the detector is for use by severely

disabled people. Someone who has severely limited movement may be able to use the detector to trigger a call for help by gazing at an alarm device for longer than a set time. Or it could be used as a binary switch to turn on or off a device such as a lock to open or close a door. If the detector was compromised, this ability is removed from the user.

Security should be considered at the design stage of any system. It should be included in the prototype and not “bolted on” as an afterthought.

The following is a list of possible threats. They are analysed by probability and cost, and then where possible, recommendations made that should be considered in any future prototype or commercial system:

Power Failure

The system has no internal battery, so if the power is lost, the system cannot signal a gaze detection event to the client system. If this was a “help” call system, then an intruder could use this attack to disable the call device.

Risk Probability: High. Risk Cost: High

Recommendations: Include a battery backup in the production model and include a “heart-beat” to signal to the client that the detector is still operational.

Occlusion

The system must be able to “see” the person to detect gaze. If the person is not in line of sight of the detector it will not be able to signal a gaze detection event to the client system. This is not something the detector can detect in that for instance, a user could be on holiday. However, if the system does not report a gaze within a certain timeframe, it should attempt to validate its own operation.

Risk Probability: High. Risk Cost: Low

Blinding

If somehow a reflector was set up correctly it could "blind" the system by tricking it into thinking it had detected a continuous gaze. This could be deliberate or an accident. Risk Probability: Low. Risk Cost: High

Recommendations: Including a sanity check in the software which can detect for an overly long gaze detection event, and for some form of notification to the client that the detector can no longer give reliable information. The detector would need to be reset to when the situation is remedied.

Physical access to the device

A malicious user could reprogram the 8051 to for example, change the threshold at which the gaze is detected. The technical requirements for this to happen are fairly advanced and the probability quite low.

Risk Probability: Low. Risk Cost: High

Recommendations: Ensure the detector is installed in a secure enclosure to prevent tampering. It is not necessary or cost-effective to have a highly secure, tamper-resistant enclosure. The following are the possible events that can occur in the system, and are classified into "correct" and "incorrect" events and can be used in evaluating the operation of the detector:

Correct Events

1. Signalling a gaze detection event when there is one to signal
2. Not signaling a gaze detection event when there is none to signal
3. Signalling inability to detect gaze correctly
4. Signalling heartbeat when system operating correctly

Incorrect Events

1. Signalling a gaze detection event when there is none to signal
2. Not signaling a gaze detection event when there is one to signal
3. Signalling inability to detect gaze incorrectly
4. Signalling heartbeat when system not operating correctly
5. Not signaling heartbeat when system operating correctly

Chapter 6

Conclusions

6.1 Summary

This project discussed the problem of context in a ubicom environment. It discussed the motivation for the prototype and introduced the concept of treating interaction in a ubicom environment as a multi-party conversation. It introduced eye-gaze detection as a useful interface for ubicom devices and covered the current state-of-the-art in gaze tracking and gaze location and focused on specific work relevant to the project. The design of the prototype was presented discussing relevant topics as such as laser technology, laser safety, red-eye and retroreflectors. The implementation of the prototype was described including the hardware acquisition and installation and the software both supplied and written. Finally the project was evaluated with reference to the original project objectives and security considerations.

6.2 Further work

This project has proven the concept of detecting reflected laser light from a simulated eye. However, further work needs to be done to optimise the prototype. Once this is complete, there is further work in preparing the product for potential commercialisation. Finally,

different application areas need to be investigated to see where the detector could be useful.

6.2.1 Optimising the Prototype

Work to refine and optimise the prototype include:

Determine minimum power laser.

Determining the minimum power laser needed to operate using visible light, and then experiment with infra-red lasers to see if eye-safe levels can be achieved. Even if the laser light is invisible, if the scan rate of the laser across the eye is fast enough that not enough energy is transmitted, eye safety can be achieved.

Improve light sensor.

There is work to investigate the best type of light sensor to use. This involves selecting a sensor based on tradeoffs between rise time and surface area, for example.

Calculate optimum scanning rate

Slower scanning rates would allow the laser illuminate the sensor for longer and perhaps increase the returned data values.

Use filters.

Use filters to increase the detecting ability of the photodiode. This would effectively exclude ambient light and reduce noise and error.

Improve model of the eye.

For testing, it would be necessary to model the eye more accurately and to record both the energy levels the eye is exposed to and hence the energy lost in reflection.

6.2.2 Towards Commercialisation

Having refined and optimised the prototype, work needs to be done to investigate how the prototype could be developed into a commercial product.

Test prototype with application

Linking the prototype to some application that uses voice recognition or gesture recognition to prove the usefulness of the interface.

Implement low power options.

Testing with low power operating modes and confirming power consumption levels

Implement a cheaper scanning solution.

Finding and testing a cheaper alternative to the scanner used. Investigate using closed-loop galvanometer technology. Investigate a possible solution using few or no moving parts.

Miniaturisation.

Investigating ways to miniaturise the entire prototype.

Identify application areas

Identify further potential areas where the detector could be used.

6.2.3 Disabled and Blind Assistance

This project also has an application for blind users. Most people registered as blind are not completely blind and have some use of their eyes. This is important not because of the ability see shapes, but because they have control over their eye muscles. Blind people have very good spatial acuity - they remember where things are. The detector could be configured to emit an audible signal when it detects gaze. This could be different for each detector device. Therefore a blind person could enter a room and “look” around to “see” what is in the room. Identifiable and standard audible signals could identify different things in the room. What’s more these things could be moved and still detectable to the user. There would be much work involved to prove that this concept is valid, but the addition of an audible signal to the prototype would not take much work. The detector could also be of use to those who have limited movement ability, but do have good eye control. Gazing at devices could allow for interaction such as turning things on or off or calling for help. Gazing at a device could

enable it to “call out” its options and then by looking away trigger that option.

6.3 Conclusion

The project has shown the idea of detecting a users gaze using a laser is possible. It has also shown how this one single action could enhance a ubicom environment. It has shown that using eye-gaze is a natural and ordinary thing for a person to do. It has also shown potential further work that could have valid benefits both academically and commercially.

Bibliography

- [1] *Part 2:1 Lasers. Safety in Universities: Notes of Guidance.* Committee of Vice-Chancellors and Principals (UK), 1992.
- [2] *Radiation Safety of Laser Products, Part 1.* National Standards Authority of Ireland, 1994. IS/EN 60825-1:1994.
- [3] Edmund optics. <http://www.edmundoptics.com/>, 2006.
- [4] Radionics. <http://www.radionics.ie/>, 2006.
- [5] Rita51, 8051 embedded controller for the internet. <http://www.rigelcorp.com/rita51j.htm>, May 2006.
- [6] D Beach, A Shotwell, and P Essue. *Applications of Lasers and Laser Systems.* PTR Prentice Hall, New Jersey, 1993.
- [7] Louise Bradley. *Laser Safety Training Manual.* Dept of Physics, Trinity College Dublin, October 2003.
- [8] Guanling Chen and David Kotz. A survey of context-aware mobile computing research. Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College, November 2000.
- [9] JDS Uniphase Corporation. Web site. <http://www.jdsu.com>, May.

- [10] Analog Devices. Ad7718 evaluation board. <http://www.analog.com/>.
- [11] Analog Devices. Ad7718 data sheet. <http://www.analog.com/>, 2006.
- [12] Paul Dourish. What we talk about when we talk about context. *Personal Ubiquitous Comput.*, 8(1):19–30, 2004.
- [13] Samuel M Goldwasser. Laser faq. October 2005.
- [14] J Goodwin and J Heritage. Conversational analysis. *Annual Review of Anthropology*, 19(3):283–307, 1990.
- [15] Jochen Heinzmann and Alexander Zelinsky. 3-d facial pose and gaze point estimation using a robust real-time tracking paradigm. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*, pages 142–147, 1998.
- [16] James E Hendrix. *The Small-C Handbook*. Reston, 1984.
- [17] D G Hunter, S N Patel, and D L Guyton. Automated detection of foveal fixation by use of retinal birefringence scanning. *Applied Optics*, 38(7), 1999.
- [18] K Irie, B A Wilson, and R D Jones. A laser-based eye-tracking system. *Behaviour Research, Instruments and Computers*, 34(4):561–572, 2002.
- [19] Douglas W. Jones. Punched cards - a brief illustrated technical history. <http://www.cs.uiowa.edu/~jones/cards/history.html>, 2005.
- [20] D Kalinsky and R Kalinsky. Introduction to serial peripheral interface. <http://embeddeddsp.embedded.com/shared/printableArticle.jhtml?articleID=9900483>.
- [21] Lincoln Laser. Sos-bb10 data sheet. <http://www.lincolnlaser.com/products/offtheshelf.cfm>, 2006.

- [22] J Magee, M Scott, B Waber, and M Betke. Eyekeys: A real-time vision interface based on gaze detection from a low-grade video camera. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW '04)*. IEEE, 2004.
- [23] T. H. Maiman. Stimulated optical radiation in ruby. *Nature*, 187(4736):493–4, Aug 1960.
- [24] Y Matsumoto, J Ido, K Takemura, M Koeda, and T Ogasawara. Portable facial information measurement system and its application to human modelling and human interfaces. In *Proceedings of the Sixth IEEE International Conference on Automatic Face and Gesture Recognition (FGR '04)*, 2004.
- [25] Joe McKnight and Eoin Meehan. Gesture recognition pilot project. 2005.
- [26] Carl Mungan. The cat’s eye retroreflector. <http://www.usna.edu/Users/physics/mungan/Scholarship/Retroreflector.pdf>, 2001.
- [27] H Nyquist. Regeneration theory. *Bell Systems Technical Journal*, 111(3):126–147, 1932.
- [28] Dept of Teaching Media. *Laser Safety in Higher Education (Video)*. University of Southampton, UK.
- [29] Kang Ryoung Park. Vision-based facial and eye gaze tracking system. *Lecture Notes in Computer Science*, 3238(4):453–465, 2002.
- [30] Gunter Scholz. Question and answer corner. *Phys13News*, (102):12, Spring 2002.
- [31] Jeffrey S. Shell, Roel Vertegaal, Daniel Cheng, Alexander W. Skaburskis, Changuk Sohn, A. James Stewart, Omar Aoudeh, and Connor Dickie. Ecslasses and eyepliances: using attention to open sociable windows of interaction. In *ETRA '04: Proceedings of the 2004 symposium on Eye tracking research and applications*, pages 93–100, New York, NY, USA, 2004. ACM Press.

- [32] John D. Smith, Roel Vertegaal, and Changuk Sohn. Viewpointer: lightweight calibration-free eye tracking for ubiquitous handsfree deixis. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 53–61, New York, NY, USA, 2005. ACM Press.
- [33] K Takemura, H Minamide, Y Matsumoto, and T Ogasawara. What you look at is what you control. <http://robotics.naist.jp/kenta-ta/paper/TEXCRA2004.pdf>, 2004.
- [34] Ashley George Taylor. Topic report: Wimp interfaces. 1997.
- [35] Global Laser Technology. V3vll data sheet (march 2006). <http://www.global-lasertech.co.uk/vector/pdf/laserlyte.pdf>.
- [36] A M Turing. Computing machinery and intelligence. *Mind*, (59):433–460, 1950.
- [37] Roel Vertegaal, Robert Slagter, Gerrit van der Veer, and Anton Nijholt. Eye gaze patterns in conversations: there is more to conversational agents than meets the eyes. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 301–308, New York, NY, USA, 2001. ACM Press.
- [38] Roel Vertegaal, Robert Slagter, Gerrit van der Veer, and Anton Nijholt. Eye gaze patterns in conversations: there is more to conversational agents than meets the eyes. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 301–308, New York, NY, USA, 2001. ACM Press.
- [39] M Weiser. The computer for the 21st century. *Scientific American*, 265(3):94–104, 1991.
- [40] Z Zhu and Q Ji. Eye and gaze tracking for interactive display. *Machine Vision and Applications*, 15:139–148, 2004.