


HAPTIC SIMULATION FOR PSYCHOPHYSICAL INVESTIGATIONS

-

by **Andrew John Hardwick**

of BT Laboratories,
Martlesham Heath



A dissertation
submitted to the University of London for the
Degree of MSc in Telecom Engineering.

1. CONTENTS

1.	Contents	1
2.	Summary	3
3.	Declaration	4
4.	Introduction	5
4.1.	Overview	5
4.2.	Benefit to BT	5
4.3.	Other Formalities	5
5.	Literature Review & Background	8
5.1.	Nomenclature	8
5.2.	Mechanics vs Mathematics vs Psychology	8
5.3.	Applications	9
5.4.	Mechanics	13
5.5.	Mathematics	15
5.6.	Psychology	17
5.7.	Devices Tested	18
5.8.	Unanswered Questions	23
6.	Design of User Experiments	24
6.1.	Informal Experiments	24
6.2.	Solid Object Experiments	24
6.3.	Texture Experiments	28
6.4.	Apparatus Requirements	30
7.	Design & Build of Equipment	31
7.1.	Selecting Hardware	31
7.2.	Interfacing Software to Hardware	31
7.3.	Simulating Solids	32
7.4.	Simulating of Textures	36
7.5.	Timing Issues	37
7.6.	Texture Experiment Application	38
7.7.	Solid Objects Experiment Application	40
7.8.	Demo' Applications	41
8.	Results of User Tests	43
8.1.	Informal Experiments	43
8.2.	Notes concerning the Statistical Analyses	49
8.3.	Solid Object Experiments	50
8.4.	Texture Experiments	55
9.	Discussion of Results	58
9.1.	General Observations	58

9.2. Tardis Effect	58
9.3. The Physical Cause of Roughness	60
9.4. Haptic Gamma Correction	62
10. Recommendations	64
10.1. BT Haptic Recommendations	64
10.2. BT Other Recommendations	64
10.3. Further Research	65
11. Conclusions	66
12. References	67
13. Appendices	75
13.1. Appendix A: Kinematics Equations for IE 3000	75
13.2. Appendix B: Class structure	77
13.3. Appendix C: Impulse Engine 3000 Design Faults	78
13.4. Appendix D: Haptic Networking Issues	79
13.5. Appendix E: Acronyms	83
13.6. Appendix F: Publications from this Work	83
14. Acknowledgements	87

2. SUMMARY

The sense of touch is becoming the third element of computer multimedia following sight & sound. This is especially useful for blind users who are being excluded by the growth of naïvely designed popular user computer desktop & World Wide Web displays with purely graphical controls but is also helpful to sighted users for whom it both aids normal use and enables totally new applications. However, it will have implications for telephone companies providing computer networks. Even the fundamentals of what needs to be stored, transmitted & output to cope with feelable computer systems are currently unknown. Questions to be answered include: “What exactly is it that gives the sensation of roughness?”; “Are there equivalents of visual illusions that will need compensation?”; & “Do there need to be settings that are adjustable between different users?”. This dissertation reports research that addresses these & other questions.

The design & build of an experimental system is described including detail of the algorithms used to simulate objects & textures. The methods & results of the psychophysical experiments for which it was used are summarised. The fundamental basis of the feeling of roughness is traced to the magnitude of the stick-slip friction as the surface is traversed. Several illusions are discovered including a strong repeatable one named ‘The Tardis Effect’ whereby objects felt from inside feel bigger than when felt from outside. A parameter akin to visual-display gamma correction is proposed as one setting that may need to be adjustable for textures to feel the same to different users. Business recommendations for BT are made.

- Version 1: 2000-07-17.
- Version 2: 2000-07-31. Corrections after viva; mostly just replacing some URL paper references with conventional print-style ones for the same papers.
- Version 3: 2024-02-07. Redacted some now-obsolete business recommendations etc. & corrected some typos.

3. DECLARATION

“This dissertation is a result of my independent investigation. The portions of the report that are indebted to other sources have been referenced in the normal manner.”

[For detail of who did what, see the final section of the Introduction]

and

“This dissertation has not already been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.”

4. INTRODUCTION

4.1. OVERVIEW

This paper describes the design of a system used to investigate the human perception of touch, some of the psychophysical experiments carried out with it and some of the resulting scientific discoveries.

4.2. BENEFIT TO BT

The work was funded by British Telecommunications plc. because of its relevance to telecommunications & computer systems in many ways. Considered especially important were:

- **Visually disabled people** are users of computers. The trend of modern computer systems from textual to graphical user interfaces is excluding them. This is disadvantageous to BT not only because of its social & legal obligations to permit disabled access but simply because disabled people are, like able-bodied people, customers. Blocking customers reduces potential revenue.
- **Network requirements** for transmitting touch, whatever it is used for in the future, need to be understood. This will reduce the risk of BT being caught unawares by forthcoming customer requirements. If, for example, previously ignored millisecond differences in latency between network providers becomes a market differentiator then it will be useful to know in sufficient time to provision for it. Understanding is also needed to develop useful Quality of Service measure.
- **Publicity** from this work is good for BT because it goes through channels that normal explicit adverts do not and its remarkably good mix of human interest, high technology & academic honesty make a seductive story for the media. (Some of publicity this project has produce is listed in Appendix F.)

4.3. OTHER FORMALITIES

4.3.1. METHOD OF ENQUIRY

The research has been carried out in the conventional manner: hypotheses were proposed; experiments were designed; apparatus was designed & built; experiments were carried out; data were analysed; the results tested the hypotheses and inspired new ones. Of course, the research went through several iterations & many mistakes & detours but it is presented below, in the recommended style for this report, as a smooth flow of methods-apparatus-results rather than in real historical order.

4.3.2. SOURCES OF INFORMATION

The information has been obtained from the research & discussion in this project or from the sources (mainly printed academic papers or the WWW) referenced.

4.3.3. ARRANGEMENT & GROUPING OF DATA

The paper is structured from the middle outwards like a sandwich filled sandwich. The meat of the project is the middle section 7, the design & build of the apparatus. This is surrounded by the psychological experiment sections, 6, the experiment design which determined the apparatus requirements, & 8, the experiments' results. This in turn is surrounded by the sections which fit the work in context, 5, which reviews the background literature, & 9, which discusses the implications of the results. The next layer is the BT justification required for the MSc, 4, the introduction, & 10, the recommendations. Then the matching précis, 2, the summary/abstract, & 11, the conclusions.

There are also appendices, acknowledgements & references at the traditional place at the end.

4.3.4. INDICATION OF STUDENT'S CONTRIBUTION

The current author wrote the whole text of this report & drew or photographed all of the figures in it (except those explicitly referenced in the captions to other sources). All the algorithmic design & programming was by the current author.

The psychological work reported here was performed in collaboration with University of Hertfordshire's Sensory Disabilities Research Unit (SDRU) which is headed by Prof. Helen Petrie. The official relationship between the SDRU & BT is that BT was sponsoring a CASE studentship for Paul Penn in the SDRU with the current author as CASE supervisor. Aside from the obvious monetary considerations there was much mutual benefit from working together. In particular, BT gets user trials carried out by the SDRU & the SDRU gets bespoke haptic systems developed for its experiments by BT. Although the primary aim of BT was to make money & that of the SDRU was to help disabled people, the science of human communication was vital to both, so the results same scientific studies were of importance to both. The degree of interaction between the partners was much higher than in many CASE PhDs.

As regards this project: the formal experiments were carried out by students

(Chetz Colwell, Paul Penn, Timo Bruns & Mark Brady) at the SDRU; the statistical analyses were performed by the SDRU students with help from lecturer Diana Kornbrot; the experimental designs were initially proposed by the SDRU students &/or their supervisor, Helen Petrie; and the designs were refined in discussion with BT team of psychologist Stephen Furner & the current author. Being the only non-psychologist involved, it was the responsibility of the current author in experimental design discussions to ensure the apparatus requirements were physically & computationally feasible. Out of the SDRU students, Chetz Colwell determined most of the basic experimental designs. It was Paul Penn & Timo Bruns carefully performing refinements of those who produced the most accurate data.

The main discoveries (in historical order):

- Alternative mental models effect: Discovered by the current author. The mixed model addition by Chetz Colwell.
- Tardis Effect: Discovered experimentally by Chetz Colwell. Formally confirmed by Timo Bruns.
- Negative Stevens' exponent: Discovered as a possibility by Chetz Colwell. Confirmed as being the dominate case by Paul Penn.
- Blind & sighted subject results being the same: Discovered by Paul Penn.
- Haptic gamma correction: Invented by the current author.
- Explanation of different Stevens' exponents: Hypothesised by the current author.

BT project/task management for haptics has involved many people - including Stephen Furner, the current author, John Seton, David Hands, Mike Hollier, Paul Barrett & Jim Rush - because its wide applicability has caused it to move into many work areas.

The project has included much additional work that is not dwelt on here, including networked haptic devices (see Appendix D) & one invention that is still confidential, of which the current author's contribution was also apparatus/algorithm design & build.

5. LITERATURE REVIEW & BACKGROUND

A detailed review of the whole science, technology & history of haptics would be book-sized so the following only covers the main aspects of haptics relevant to this project. These are: a brief glossary; applications of haptics; the mechanics & mathematics of haptic simulation; & the prior psychophysical studies into textures & shape. For broader reviews, see Srinivassan & Basdogan 1997 about simulation and Loomis & Lederman 1986 about psychophysics.

5.1. NOMENCLATURE

At this stage it is useful to define some common terms. (For acronyms see Appendix E.)

- **Tactile:** relating to the sense of touch from nerve sensors in the skin [see Sherrick & Cholewiak 1986 for detailed review].
- **Kinaesthetic:** relating to the sense of touch from sensing the position of the body's joints [see Clark & Horch 1986 for detailed review].
- **Haptic:** relating to the sense of touch in general including both tactile & kinaesthetic.
- **HCI:** Human-Computer Interaction ('CHI' in US English).
- **Psychophysics:** the science of human mental interaction with the physical environment (typically concerning unconscious low-level data-processing in the brain).
- **Virtual Reality:** an oxymoronic but fashionable term meaning computer simulated environments (particularly for HCI).

5.2. MECHANICS VS MATHEMATICS VS PSYCHOLOGY

Upon entering the field of haptic HCI, one immediately finds that researchers are essentially split into three groups. Mechanical engineering is the obvious group because physical equipment must be designed for haptic output. Mathematics is also needed because some algorithm must be used to calculate an appropriate output for each situation being simulated that is stable, convincing & efficient. Psychology, although often ignored, is vital because the whole purpose of HCI is to interact with a human. It is rare for a paper to cover more than one of these, others being assumed

obvious, but the work reported here needed all three: suitable hardware had to be selected; algorithms had to be created for texture & solid simulation; and psychological experiments had to be designed and carried out.

The three will be reviewed in order but firstly blind access & other uses of haptics will be covered to justify the research.

5.3. APPLICATIONS

5.3.1. COMPUTER & WWW ACCESS FOR VISUALLY DISABLED USERS

Computer access by for blind users is actually being hampered by the spread of 'user-friendly' the graphical user interface (GUI). The text of a purely alphanumeric display can be automatically converted to speech output or Braille. Only a few aspects, such as tables and ASCII-art have significant information in the visually presented spatial arrangement on-screen. However a Windows-Icons-Menus-Pointer (WIMP) system is heavily visually biased making use of the ability quickly scan a scene with structural clues by sight.

A particularly heated issue is the World Wide Web (WWW). This is very popular amongst visually impaired computer users as a source of information. The language of WWW pages, HTML, is ideal for this. It was designed as pure text plus structural indications to be used to guide formatting at the terminal. This is so that if a terminal cannot support fancy formatting or the network is too slow for graphics, the output gracefully degrades into a still usable, just less pretty, form. A classical example of a modern restricted terminal is a palm-top computer linked to a mobile telephone. Not only is HTML good when the hardware is restricted but also when the user's senses are restricted. It can be displayed as speech or Braille and automatically restructured (such as bringing the links to the top of the page) to optimise for the serial nature of the output.

Unfortunately now that the WWW has become so popular with the general public, this good design is being compromised through a mix of greed and incompetence. Information that could be efficient versatile HTML becomes text drawn as pictures, explicit font & colour commands, framesets, Javascript, Java, Shockwave, etc.. The greed came in with commercialisation. To attract gullible customers, commercial sites sacrificed accessibility for window dressing not realising, or not caring, that they have lost two whole market sectors - blind & browser-limited customers. Many of today's WWW authors are no longer knowledgeable about the fundamentals of HTML & how to structure documents for

the medium. Instead they just assume that what looks okay to themselves (or their bosses) on their own particular computers will look the same anywhere & are even oblivious to possibilities of different screen widths or colour blindness, let alone text displays and full blindness. Educating WWW authors is probably fruitless. For example I got *[redacted business information]* fixed, by direct contact to the page designer, but only the front page and then only lasting until the next change in advertising design. Without accessibility legislation (which is difficult to implement internationally), one may have to rely on technology to reinterpret the graphical pages as well as the text in the browser. This is where haptics may come in*.

Maybe haptic displays can do for the GUI what Braille output did for text. They have already been used to present computer output that is intrinsically pictorial such as 2d graphs & 3d data visualisation spaces [Brewster & Pengelly 2000]. Haptics could be used as 'assistive technology' add-ons helping blind people with information in otherwise inaccessible media [Petrie 1997], 'adaptive technology' changing otherwise inaccessible systems [Petrie 1997], or in 'design for all' which is the careful design of systems so that virtually everyone, disabled or not, can use them [Ekberg 2000]. This latter design philosophy is particularly apt for the telecommunications industry because it was in developing a hearing aid that Alexander Graham Bell invented the telephone.

Disabled accessibility is very important to BT for 2 commercial reasons. Firstly, disabled people are customers just like able-bodied people and not supplying products & services to any market segment reduces potential revenues. The disabled are not a minor segment. They comprise about 10% [Petrie 1997] of the population & the percentage is rising with increasing numbers of old-folk. Secondly, companies have legal obligations to allow disabled access; making public computer services inaccessible to blind users is comparable to removing wheelchair ramps from public buildings. Ignoring these obligations can result in legal action. For example, in the more litigious USA, the National Foundation for the Blind has already sued AOL, the popular internet service provider owned by WorldCom, for inaccessibility [Vaas 2000].

* However haptic, audio or whatever versions of graphical WWW pages are a poor substitute for the sensible practical approach of having WWW pages designed well in the first place so that they are useable by anybody on any medium. What is really needed is to educate WWW authors & their employers. I could write several essays on the practicalities, benefits, ethics, etc. of accessible WWW sites but this dissertation is on haptics so it is the haptic option that I will concentrate on.

Worse may still be to come. There have been several proposals (e.g. VRML) to use 3d 'virtual reality' GUI instead of 2d WIMP. Fortunately the expense of satisfactory stereoscopic displays, the crudeness of general purpose 3d scenes and the slowness of use are currently deterring this but it is an ominous future possibility.

5.3.2. IMPROVING GENERAL HCI

The tools that people normally work with have haptic feedback. One does not need to be visually or aurally waiting for a response when moving ones hands across a workbench to check that one has reached a tool before grasping it, unlike when moving a mouse to click an on-screen button. For years, both the general public and professional typists have preferred moving key computer keyboards to flat unresponsive ones (as on the ZX81). Adding haptic feedback to computer input devices could allow computer workers to work faster, easier & more accurately. The little research which has been done on this seems to confirm this expectation:

- **Haptech pantograph:** Tested for use in free-fall for the Canadian space agency, it was significantly better than a conventional trackball (although ground-tests implied most of the improvement could have been from the pantograph not the force-feedback) [Payette *et al* 1996].
- **Immersion mouse:** When quickly moving between targets on screen, speed improved by 61% & errors reduced by 70% with force-feedback [Hasser *et al* 1998].
- **Phillips trackball:** Force-feedback improved efficiency, especially when the haptics gave not just information about what was there but guided the user towards outcomes [Engel 1994].

(However, note that all the above studies were performed by the manufacturers who had reason to promote their devices.)

One has to be careful when adding haptics to a visual output though because haptic & visual illusions differ. These can conflict resulting in combined performance that is little better than with vision alone if underlying principles are not understood [Fukui & Shimojo 1992]. This highlights the importance of understanding the psychophysics not just the hardware.

5.3.3. OTHER USES OF HAPTIC HCI

There are many other obvious uses of haptics including:

- **Entertainment:** Computer games are the initial mass-market for cheap haptic devices. Players are not fussy about lack of realism in the simulation. Haptics is more of a gimmick increasing the feeling of immersion than providing useful information (our informal tests with a Sidewinder Force-feedback Pro showed that it was easier to play with the force-feedback removed). Conveniently many game controllers are already shaped like controls that would naturally vibrate (e.g. joysticks & steering wheels) so simple wobbles suffice without having to simulate the much more complex direct feeling of objects.
- **Computer Aided Design:** A mechanical engineer using a CAD package with added haptics could feel the stiffness of the drafted item as well as seeing it. This is likely to be an early market for mid-price devices.
- **Data visualisation:** There are certain plots which may be easier & more natural to feel explored 3d than see in 2d with occlusion. For example 3d vector-fields like electromagnetic fields, fluid flows & stress distributions [Infed 1999].
- **Manual skill training:** Using haptics to extend computer-based training to manual skills sounds reasonable but initial results are disappointing. A study where a remote instructor haptically guided the trainee's hand showed increased learning speed but no improvement in long-term retention and, after teaching people how to reduce unwanted swinging of crane loads, haptics was found to have no benefit & all the subjects who tried the haptics said verbal instruction would have been better [Gillespie *et al* 1998].
- **Medical training:** Humans, or animals, to practice on are in short supply, or are expensive, so haptic simulations are particularly desired. A classic example was the simulation of the feeling of a horse's ovaries by a finger inserted in her vagina [Brewster *et al* 2000] where the traditional teaching method can vex the mare felt by a whole class of trainee vets. In human medicine, endoscopic ('keyhole') surgery [Cao *et al* 1996] is a common haptically simulated training application both because it a difficult skill to learn & because tools are relatively easy to simulate essentially being sticks pivoting through holes.

- **Electronic music:** Acoustic musical instruments give immediate haptic feedback from the sound vibration which is missing from many electronic ones unless simulated [Bongers 1997].
- **Rehabilitation:** Simple haptic stimulation can be used for controlled exercise of joints & muscles after injury [Gione 1999].
- **Teleoperation:** Controlling mechanical machines remotely [Hannaford & Venema 1995].
- **Telemedicine:** Always a popular teleoperation example in haptic hype is possibility of carrying out surgery remotely. However, the quality & reliability needed for this and the uncertain legal liability for failures in the system make this unlikely in the near-term other than maybe for battle-field use.

5.4. MECHANICS

5.4.1. POSSIBLE PRINCIPLES

Mechanical haptic simulation can be based on any of several different principles. Non-mechanic stimulation, e.g. by heat [Ottensmeyer 1997] or mild electric shocks, has been used but is not common. Most devices use 'force-feedback'. It is so prevalent that 'force-feedback' is often erroneously used as a synonym for 'haptic'. Alternative principles include:

- **Force-feedback:** The user moves an easily-moveable end-effector. The force that the user would feel if the simulated objects were real is calculated from the position & this force is mechanically applied to the end-effector.
- **Displacement-feedback:** The reverse of force-feedback. A difficult-to-move end-effector is pushed by the user. The position that it would move to if the simulated objects were real is calculated from the force & the end-effector is mechanically moved to that position [Hannaford & Venema 1995].
- **Shape fitting:** A complex mechanism physically changes to the shape being simulated & is then felt by the user. Typically the arrangement is just a 2d array of actuators than can be extended to simulate a bumpy surface.
- **Contact signal:** Some arbitrary haptic signal is used to indicate when an object is contacted. For example a vibration might be applied to a finger tip when it reaches a simulated surface [Murray 1999].

- **Tool simulation:** Instead of simulating contact with a surface, the hardware might be specifically designed to mimic a tool like a joystick or rotary switch. This is easier because it is one mechanism simulating another mechanism but is less versatile.

The reason for the popularity of force-feedback over displacement-feedback is that its no-power loose state corresponds to the free space which makes up the majority of the reachable volume in typical simulations. Shape-fitting is mechanically complicated & expensive needing many actuators. For example the box of actuators to drive a 20 x 20 dot array sized 1 cm² was about 1 m³ [Pawluk *et al* 1998]. The only mass-marketed devices using dot arrays are Braille displays & even those could be considered tool simulations.

5.4.2. POSSIBLE ACTUATORS

There are many possible actuators that can be used for haptic outputs:

- **Electric motors:** Virtually all force-feedback & displacement-feedback devices use motors. These are convenient & controllable by force or position but are heavy & bulky so that, when more than few degrees of freedom need actuators, the result is a mass of cables leading to the motors which increases friction & inertia.
- **Pneumatic bellows:** These have often been used in experimental haptic devices from university robotics departments [e.g. Pere *et al* 1996] but the resolution & speed achievable is not generally good.
- **Solenoids:** These are typical linear actuators for big dot array system.
- **Piezoelectrics:** Bi-piezoelectric strips are used for Braille cells. These are relatively bulky compared to the dots because a lever is needed to magnify the deflection but Braille letters only need to be 3 or 4 dots high & the levers can be hidden under a keyboard.
- **Memory metals:** These are wires which contract when heated by internal electric currents. They were fashionable in the early 1990s but cooling requirements limit response times to over 1/3 s [Taylor *et al* 1997b].
- **Magnetic levitation:** This has been used for a floating handpiece with 6 degrees of freedom (translation & rotation) but it only had $\pm 7^\circ$ & ± 13 mm range [Berkelman *et al* 1996] & has not progressed [Berkelman & Hollis

1998].

- **Electrorheological fluids:** Only in concept demonstrators [Taylor *et al* 1997a].
- **Ultrasonic standing waves:** These can simulate different friction coefficients [Nara *et al* 1999].
- **Silicon microengineering:** Although it was a hopeful solution for fine feelable dot arrays, in practice silicon was too brittle and moveable parts broke off when felt [Waller 1995].

5.4.3. POSSIBLE SHAPES

Force-feedback outputs can come in many shapes for different applications:

- **Imitating direct touching:** The handpiece or finger-piece moves as if the hand or finger is directly touching the simulated objects. This is often assumed as an ideal for immersive interaction in simulated environments but is slightly unreal because what the finger is actually touching moves with it so the expected tactile sensations are removed, leaving only those from secondary kinaesthetic effects.
- **Mice & trackballs:** These mimic conventional WIMP pointing devices.
- **Steering wheels, joysticks, guns:** These mimic conventional computer game controls.
- **Gloves:** These are intuitively attractive the difficulty of fitting in actuators for the 20 degrees of freedom of the fingers makes them expensive [e.g. the Cybergrasp costs £35k] & rare compared to glove sensors which have available from several suppliers for years [Sturman & Zelter 1994]
- **Free-floating:** These are not practical with current technology but cable-suspended end-effectors [WilliamnsII 1998] come close.

5.5. MATHEMATICS

Much of the literature anent haptic simulation algorithms (e.g. Miller 1999) is on the detail of how to mathematically ensure stability of the haptic image and how best to optimise collision detection (e.g. Pai & Reissell 1996, Cai *et al* 1999). This is typically a theoretical exercise & often never actually implemented. More relevant

are those practical algorithms made *ad hoc* in the process of the creating real haptic displays. However, these are so numerous with so many options that a full survey would be inappropriate here. Only a brief overview of some alternatives for force-feedback follow. The particular algorithms used in this study are described in detail in the Design & Build section.

The general principle of simulating solids is to deter the users from entering objects with appropriate forces. This requires a mathematical representation of the simulated scene. There are two main approaches to this: representing the surfaces as algebraic forms mapped out from splines, planes & points (e.g. Thompson.II & Cohen 1999; or building the scene from groups of basic shapes like spheres & cubes. The later is used both in this research, because of an initial requirement to be display VRML, & by the Phantom's GHOST API [Massie 1996]. Of course, the approaches are not mutually exclusive; algebraic surfaces can be built up in groups or the set of basic object can include a general-purpose polyhedron.

There are many ways of simulating textures on surfaces. Some use a visual texture bitmap as a source (e.g. Ho *et al* 1999), some use a deterministic mathematical model as a surface displacement map (e.g. Hardwick 2000), & some use a statistical mathematical model (e.g. Green & Salisbury 1997). The bumps of a texture can be simulated in 2d using only in the in-plane component of reaction force (e.g. Minsky & Lederman 1996), 3d allowing users to also break contact & skim across the bumps naturally (e.g. Hardwick 2000), or with the actual surface boundary simulated in 2d or simplified 3d combined with reaction forces adjusted to mimic full 3d [Morgenbesser & Srinivasan 1996].

Most systems simply rely on the programmer to define the scene but a few try to record reality by scanning the handpiece across a real environment & building up a model of what it encounters. This creates additional decision of how to store the recording. Options include spatial stiffness records [MacLean 1996], texture force Fourier series [Wall & Harwin 1999], & wavelets [Miller & Colgate 1998].

A decision must be made as to how to relate the physical handpiece position to position in the simulated space. The most direct representation is to consider the two positions the same (with translation & scaling if desired). This is simple to calculate, intuitive for direct touch & good for examining objects by feel but the simulated workspace is limited by the physical one. Alternatively, the handpiece position could determine the velocity of the simulated position as with a normal joystick. This allows unlimited spatial volume but makes the concept of actually feeling objects rather unnatural. Combinations are possible such as using the velocity method to

navigate to an object & the position to feel it or using just the position method but moving the simulated scene by the velocity method when the edge of the workspace is reached (like automatic scrolling in GUI drag & drop). In practice, the position method dominates except in games where the haptic effects are principally gimmicks.

5.6. PSYCHOLOGY

There has been very little psychological research performed on haptic simulations, or indeed the haptic sense at all, compared to that on visual simulations & the visual sense. Even such obvious experiments as testing if people perform tactile tasks worse wearing thicker gloves (conclusion: they do) were only performed recently [Shibita & Howe 1999].

Methodologies appropriate to visual sense investigations are not necessarily applicable to haptic ones because of the difference between the senses. Whereas vision is a passive sense, the haptic sense works very poorly at identifying objects if they are simply pressed to the hand or even if the hand is guided in exploration compared to when the hand is allowed to explore freely. This has given haptics an undue reputation as a poor sense in the past [Lederman & Klatzky 1987]. Moreover, haptic recognition is a separate cognitive process after sensing [Revesz 1950], even for basic factors like shapes of rectangles [Appelle *et al* 1980], unlike basic visual recognition which is immediately performed by from low-level processing. There have been some multimodal studies where vision and haptics have been used together; Hendriz *et al* 1999 concluded that, in comparing materials, haptics was important in matching like materials whereas vision was better at rating them individually.

The area of haptics that has received the most psychophysical attention is the perception of roughness. The pioneering work of Stevens & Harris 1962 involved subjects assigning subjective roughness values to samples of sandpaper. They discovered that the roughness was proportional to the grit number of the sandpaper raised to a constant power, β .

$$(\textit{Perceived roughness}) \propto (\textit{Grit number})^\beta \quad \{0\}$$

They also found that asking for 'smoothness' instead of 'roughness' gave inversely proportional values & that asking for a roughness to be represented as loudness of sound gave a similar law with an exponent equal to the sum of the roughness to number & number to loudness exponents. That study greatly influenced subsequent experimenters who typically followed the same process of asking subjects for

roughness numbers & fitting Equation 0 to them. Equation 0 & the power β are often called 'Stevens' Law' & the 'Stevens' exponent' respectively and are used indiscriminately for any texture specifying physical property (e.g. groove width, bump spacing or force amplitude) not just grit number.

The main researcher currently working on roughness perception is Susan Lederman who has been gathering data on it for 3 decades [Lederman & Taylor 1972, Lederman 1974, Lederman 1981, Lederman & Klatzky 1998, Klatzky *et al* 1999] using plates with parallel grooves or spaced bumps and with variations including: moving the plate instead of the finger; constraining the motion speed; feeling with probes of various widths; and constraining the applied force.

There have been 2 groups besides ourselves publishing psychophysics from simulated textures. The first was by a student of Lederman's at MIT, Margaret Minsky, using a force-feedback joystick to simulate grooves, grids & Perlin patterns in 2d with the 'Sandpaper System' software written by her husband, Oliver Steele [Minsky 1995]. She did both the standard Stevens' law fits & grouping studies applying cluster analysis to how people sorted textures. The other was Gunnar Jansson's group at Uppsala University [Jansson *et al* 1999] using a Phantom to simulate sandpaper with algorithms from MIT [Green & Salisbury 1997].

The Stevens' exponents found in these studies differed (not be unexpected given the variety of physical parameters used interchangeably to characterise the textures) but in every case it was the coarser sandpaper, wider grooves or further spaced bumps that felt rougher.

5.7. DEVICES TESTED

There is not space here to review all the haptic devices ever invented, most of which are one-off research devices anyway, but the ones we have made or acquired are briefly covered below.

5.7.1. TINI TACTOR

This was a memory metal actuated lever about 2 cm long with about 1 mm tip motion. The output was a single dot designed to touch a finger pad. The advertising exaggerated its applications even suggesting uses in remote surgery but in practice it was a pathetic 1d haptic dot with slow (~1 Hz max) response & only two states (up or down).

5.7.2. BRAILLE CELL MOUSE

Like many people new to haptics, we initially thought of dot-arrays. It was obvious that a dot array large enough to represent a GUI was impractical with current technology so I constructed a 4 x 4 dot array with 2 mm inter-dot spacing from two Tiedman piezoelectric F.S. Braille cells (Figure 1) attached to a computer mouse (Figure 2) to emulate a larger array. The system could simulate a GUI (real GUIs required finer dot spacing) and reproduce the on-screen pattern on the array via a PCIB40 digital output ISA card.




Figure 1: Tiedman Braille cell longitudinal cross-section with two dots raised & two dots lowered.




Figure 2: Early illustration of dot-array equipped mouse.

Three modes of presentation were made: direct representation of the GUI dot-pattern as a tactile dot-pattern; contour-map representation; & a fixed icon set (a noughts & crosses game). The problem of static patterns pressed to the skin soon becoming unnoticeable was solved by scrolling or blinking the pattern on the array.

The results were not encouraging. Only simple bold lines could be felt in the direct representation because of the low resolution, only a few distinct icons could be represented for the same reason, & the contour-map was useless because people did not intuitively transpose the artificial visual skill of contour visualisation to haptics.

The low resolution of the Braille cells, although mainly for mechanical reasons, was fine for their intended purpose because dots in Braille are intended to be feelable

individually not as a texture and by rubbing a finger them across not by them being raised up to the fingers. Others have put a similar 4 x 4 array on a Hapttech force-feedback mouse for Braille use. The result was called Pantobraille [Ramstein 1996].

5.7.3. OPTICON

The Opticon [Ikei *et al* 1999] has a pins array that vibrate under a finger tip with the amplitude of each being related to the brightness on-screen with the whole array representing about a cursor-size area. It is useable by people who don't know Braille but, for those who do know Braille is much slower because the letter shapes are not optimised for touch unlike Braille. It is also tiring on the finger. The first blind person I met who had used one commented that he hated having to use it at school & had abandoned it as soon as he could.

5.7.4. IMMERSION IMPULSE ENGINE 3000

This force-feedback device (Figure 3), hereafter abbreviated to IE 3000, consisted of a metal stick that could tilt like a joystick about two perpendicular axes and slide in-and-out along its length. The stick's tip, held by the user, can thus be moved in three dimensional space with encoders & motors for each axis. It cost about £7k in 1995.




Figure 3: Drawing of an IE 3000.

In its original state, it was designed to be used in a different orientation to that shown in Figure 3 (rotated 90° anticlockwise with the rod vertical) but the weight of the motors pulled the stick distractingly to one side in that orientation. It also had a pen-like stylus attached to the rod tip by a universal joint that was so slack that it masked textures & had to be discarded. It had a nominal 13 x 23 x 23 cm³ workspace and its nominal 23 μm spatial resolution, forces up to 9 N in 5 mN steps, and 650 Hz

bandwidth were better than the far superior Phantom (see below) but these were not real values, only theoretical ones based on the highly invalid assumption of perfect engineering. This shows the need for objective measures of force-feedback performance which are currently lacking. There were many other faults, see Appendix C.

5.7.5. MICROSOFT SIDEWINDER FORCE-FEEDBACK PRO

This (Figure 4) was a gaming joystick with 2d force-feedback. It was not designed for at simulating real objects but for gimmicky haptic effects & crudely mimicking the mechanical feedback on a joystick that is being used to control a real machine. Its big backer forced its API into most new home computers by incorporating it in the Windows DirectX 5 API so it is compatible with many games. Its main advantage was that it was mass-produced and cheap (launched at £130 in 1999, currently £80). Although it was from Microsoft, it appears to be robustly constructed, unlike Microsoft's operating systems, and backward compatible, unlike Microsoft's office software.

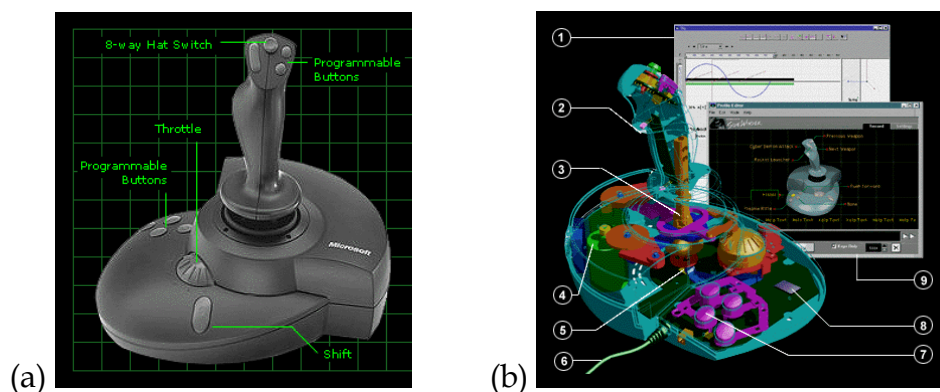


Figure 4: Sidewinder force-feedback joystick (a) normal & (b) cut-away views [pictures from Microsoft].

5.7.6. IMMERSION/LOGITECH WINGMAN FORCE-FEEDBACK MOUSE

This (Figure 5) was invented by Immersion & sold as a Logitech product. It was a mouse with 2d force-feedback from a pantograph arrangement in the mouse pad to which it is fixed (Figure 6). It is principally aimed at the gaming market by Logitech and is therefore cheap (launched this year at £80) but Immersion also promote it as a desktop GUI enhancement with software to give bump &/or friction markers to WIMP components such as buttons, window frames & menu items.



Figure 5: A Wingman Mouse.



Figure 6: Inside a Wingman, underside view.

Immersion have bought out Haptech who produced its only serious rival, the mouseCAT (Figure 7).



Figure 7: The mouseCAT [extracted from an advertising picture from Haptech].

5.7.7. SENSABLE PHANTOM 1.0

The Phantom (Figure 8) has been best readily commercially available force-feedback device for several years. It was a 3d crane-like system with interchangeable

stylus & thimble handpieces for tool use or direct contact principle simulations. The hardware was extremely well built with a nominal $13 \times 18 \times 25 \text{ cm}^3$ workspace, 0.03 mm spatial resolution, forces up to 1.4 N continuous or 8.5 N instantaneous & 1 kHz bandwidth yet only 0.04 N backdrive friction & 75 g inertia. The API, GHOST, is so comprehensive that it includes not just algorithms for reaction forces & friction taking into account entry points but a full 3d modelling language with call-backs for linking to motion algorithms & graphical simulations.

It cost £11k at time of purchase but over £20k at the time the inferior IE 3000 was purchased.



Figure 8: Phantom 1.0.

5.8. UNANSWERED QUESTIONS

Before haptics can be used for accurate output to blind people or compactly sent across a network to the benefit of BT, the basic psychophysics underlying it must be uncovered. What physically is it that actually makes people perceive one texture as rough & another as smooth? Do people feel shapes & sizes accurately without the visual clues? Indeed can a crude single point-of-contact simulation on a force-feedback device feel usefully close to reality at all? It is not new for psychophysical studies to be performed to understand requirements for telecommunications; for example a classic series of experiments on Morse telegraph operators were performed from 1893 to 1896 [Bryan & Harter 1898].

The experiments described below endeavoured answer these fundamental questions.

6. DESIGN OF USER EXPERIMENTS

6.1. INFORMAL EXPERIMENTS

Haptic computer interfaces are such a new field that there are still many scientific discoveries to be made from informal observations of people interacting with the system & from their comments. Indeed it is commonly found that, in user interface design, most of the important issues are discovered with the first few subjects [Virzi 1992] with the purpose of the remaining subjects being to formally check the statistical significance of those discoveries.

Of course, the first of these informal experiments consist of the system developers testing their system on themselves but it is also useful to have a less biased sample of subjects from the general population. These also have the advantage of not knowing how the system internally works or having practised on development prototypes and so are less likely to subconsciously compensate for the system's deficiencies. A convenient way to get such subjects is 'opportunity testing', simply demonstrating the system to interested visitors & at public shows.

6.2. SOLID OBJECT EXPERIMENTS

These experiments were to consist of presenting haptically simulated solid objects in order to detect the discrimination & repeatability with which people could feel them. This is fundamental to the use of haptics in displays. If people cannot distinguish objects then the display would be of little use. If different people need different simulations to perceive the same output then systems will need to be customised to users. If people's perception of sizes is not the same as in reality then that may have to be compensated for.

The factors that needed to be investigated included:

- The estimation of sizes of objects of identical shape but different size.
- The estimation of sizes of objects of identical size but different shape.
- The estimation of sizes of the same object felt from inside versus outside.
- The estimation of angles.
- The effect of varying the surface hardness.
- Whether or not there was any difference between sighted and blind people.

The collection & analysis of data for most of these studies were carried out by under- & post-graduate students at the Sensory Disabilities Research Unit at the University of Hertfordshire. Their names are listed below as the 'experimenters' in

the summary of experimental parameters that introduces the description of each experiment.

Presenting the experimental methods grouped together here in one section and presenting the results grouped in another section later on is the traditional scientific idiom but obscures the progressive design of these experiments. Each experiment was designed to confirm or further investigate the discoveries made in the previous experiments.

6.2.1. SOLID OBJECT EXPERIMENT 1

- Purpose = initial investigation.
- Hardware = Impulse Engine 3000.
- Cube widths = 18, 27, 36, 45 mm.
- Sphere diameters = 27, 36, 45 mm.
- Rotation angles = 30, 50, 70°.
- Corner angles = 18, 41°.
- Answering method = multiple choice pictures.
- Duplicate trials per subject & parameter combination = 3.
- Subjects = 9 blind + 13 sighted.
- Experimenter = Chetz Colwell.

This was an initial experiment and it was not even certain if sufficiently many subjects would be able to distinguish cubes from spheres let alone judge their sizes. Cubes and spheres were simulated at different sizes and the subjects asked to select by multiple choice from 2d pictures (formed from felt fabric for the blind subjects) the closest size picture to the 3d simulated ones they had felt. The pictures were of 10, 15, 20 & 25 mm cubes and 15, 20 & 25 mm spheres. The presentation of simulated & drawn objects to different scales was not intentional but an unfortunate calibration error that was not detected until after all the data had been collected. All simulated objects were presented in the centre of the workspace and the edges of the cubes were aligned with the Cartesian axes.

Rotated cubes (Figure 9) & sheared cubic (Figure 10) hollows of 36 mm height were used and compared to (scale) drawings of cubes rotated & sheared to the same angles.




Figure 9: Rotated simulated cubes.




Figure 10: Sheared simulated cubic hollows.

All the simulated objects had slightly roughened surfaces (sinusoidal grid with 0.74 mm period & 0.018 mm amplitude), because purely smooth simulated objects are as difficult to feel as oily real ones, and maximum surface reaction force of 8 N.

6.2.2. SOLID OBJECT EXPERIMENT 2

- Purpose = confirmation.
- Hardware = Impulse Engine 3000.
- Cube widths = 19, 31, 51 mm.
- Answering method = verbal size estimate.
- Duplicate trials per subject & parameter combination = 1.
- Subjects = 4 sighted.

The previous experiment had (see Results section) unearthed an unexpected haptic illusion but suffered from a calibration mistake. The discovery had been submitted for publication [Colwell 1998a, 1998b] but the mistake reduced our confidence in claiming the discovery. This was a quick simple test inspired by the hope that if a real effect was strong enough to be apparent despite the calibration mistake, it could be strong enough to be detected with very few trials.

For speed of set up, subjects were simply asked to estimate sizes in centimetres, millimetres or inches. Although this approach relies on an ability not universal in the general population, the subjects in this particular experiment were all engineers (the first 4 colleagues to walk into a lab at BT) who could be expected to be able to estimate sizes. For confirmation, they were asked to estimate sizes from drawings.

6.2.3. SOLID OBJECT EXPERIMENT 3

- Purpose = confirmation of previous results & hypothesis testing.
- Hardware = Impulse Engine 3000.
- Cube widths = 20, 30, 40, 50 mm.
- Sphere diameters = 20, 30, 40, 50 mm.
- Hardnesses (surface force limits) = 5, 6, 7, 8 N.
- Answering method = sliding sleeve ruler.
- Duplicate trials per subject & parameter combination = 1.
- Subjects = 7 blind + 10 sighted.
- Experimenter = Timo Bruns.

This was to a thorough & rigorous confirmation of the size perception effects found in Solid Object Experiment 1. It would also test a hypothesis that could explain the haptic illusion discovered (see Results section) by varying the hardness of the objects. Instead of repeating each trial several times varying one parameter at a time, a full mixed design was used with cube/sphere, size, inside/outside & hardness combinations giving 64 different objects per subject.




Figure 11: Experiment in progress with blind subject [photograph from Bruns 1998]. The guide dog was not part of the experiment.

To avoid the limitations of discrete multiple choice answers without requiring subjects to be able to estimate size units, a ruler with sliding sleeves was used. These could be visually or haptically adjusted to a size the perceived size of the simulated object by the subject and read by the experimenter. This also removed the difference in mode between stimulus & response which could have been a confounding factor.




Figure 12: Ruler with sliding sleeves in use [edited from a photograph from Bruns 1998].

6.2.4. SOLID OBJECT EXPERIMENT 4

- Purpose = confirm results are hardware independent.
- Hardware = Phantom (thimble & stylus).
- Cube widths = 27, 36, 45 mm.
- Sphere diameters = 27, 36, 35 mm.
- Corner angles = 18, 41, 64°.
- Answering method = sliding sleeve ruler.
- Duplicate trials per subject & parameter combination = 1.
- Subjects = 10 blind + 14 sighted.
- Experimenter = Paul Penn.

A obvious comment from referees when the results of the previous experiment were publish was that the findings might only apply to the Impulse Engine 3000. At the time of that experiment, it was financially impractical to duplicate it on different hardware so the experiment was later re-run after a Phantom was purchased.

The Phantom came with two ready-made handpieces, a thimble to simulate direct touch & a stylus to simulate touching with a tool, so the experiment was run with each of them so that that effectively 3 different hardware systems were tested. This also enabled the comparison between the effects of different handpieces without the confounding complications of totally different systems (e.g. Weisenberger *et al* 1999 not only changed handpieces but changed from 2d to 3d, changed mechanism & changed algorithm simultaneously).

For the angle estimation of cubic hollows, the subjects were asked to respond by flexing a carpenter's folding ruler to match the angle they felt.

6.3. TEXTURE EXPERIMENTS

The fundamental part of the texture experiments was the haptic display of simulated textures and the subjects assigning numbers to them representing their subjective impression of the roughness. Each sample consisted of a simulated horizontal plate which was flat except for a region 40 mm wide containing a series of grooves running front to back on the upper surface (Figure 13). The subjects felt across the grooved region from left to right once only and gave a roughness number

which was recorded. This was repeated for 10 samples with different groove widths but constant amplitude (Figure 14); the samples being presented in a random order. This was repeated several times for each subject and then across many subjects.

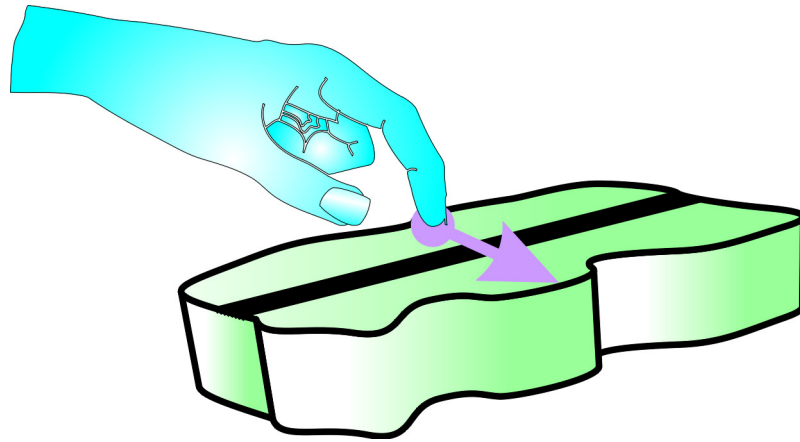


Figure 13: The simulated texture sample: an infinite horizontal plane with a band of grooves on its surface. The subject's finger moves left to right on the plane across the grooved region.




Figure 14: Definition of groove width & amplitude*.

The choice of groove width as the main parameter to vary was because it had been clearly shown to be the primary determinant of roughness in experiments on real textures. Many of the details (such as the groove widths used & the size of the textured region) were chosen to match those used by in the previous studies on real textures [Lederman 1974, Lederman 1981, Lederman & Taylor 1972] so that comparisons could be easily made.

6.3.1. TEXTURE EXPERIMENT 1

- Purpose = initial investigation.
- Hardware = Impulse Engine 3000.
- Groove widths = $0.225(n+3)$ mm, $n = 0, \dots 9$.
- Groove amplitude = 0.0625 mm.

* 'Period' & 'amplitude' have standard meanings in maths & physics but in psychology the definitions are more variable giving factor of two uncertainties. The 'amplitude' may mean the 'peak to peak amplitude' & imprecise term 'width' is used instead of 'period' even though it could equally well refer to only the negative part of a cycle rather than a whole cycle. This has caused problems in this project in the past.

- Groove shape = sinusoid.
- Duplicate trials per subject & parameter combination = 6.
- Subjects = 9 blind + 13 sighted.
- Experimenter = Chetz Colwell.

The widths & amplitude were initially chosen to match those of Lederman but a miss-calibration multiplied them by a factor of 1.8. Fortunately a constant scaling factor should not affect the resulting Stevens' Law exponent. The grooves were sinusoidal rather than rectangular like Lederman's for practical reasons. The combination of a hard point contact and sharp vertical edges required accelerations beyond the capabilities of the hardware.

6.3.2. TEXTURE EXPERIMENT 2

- Purpose = confirm results are hardware independent.
- Hardware = Phantom (thimble & stylus).
- Groove widths = $0.225(n+3)$ mm, $n = 0, \dots 9$.
- Groove amplitude = 0.1125 mm.
- Groove shape = sinusoid.
- Duplicate trials per subject & parameter combination = 6.
- Subjects = 10 blind + 14 sighted.
- Experimenter = Paul Penn.

As with Solid Object Experiment 4, this was a check that the previously found results were hardware independent. The widths & amplitude were initially chosen to match those of the previous experiment.

6.4. APPARATUS REQUIREMENTS

In order to carry out the above experiments, a haptic simulation system was needed that could:

- Simulate textures.
- Simulate 3d solid objects.
- Automate repetitive experiments & data capture where desired.
- Work with different hardware platforms.
- Be used in entertaining & informative public demonstrations.

Such a system was created and is described in the next section.

7. DESIGN & BUILD OF EQUIPMENT

This section describes creation of the experimental system from choice of hardware to algorithms for solids & textures to the experiment & demo' applications.

7.1. SELECTING HARDWARE

Whilst it would have been convenient, and potentially profitable, for us to have built a force feedback device, *[redacted business information]* limited the choice to buying a commercially available device. The choice of hardware was determined by availability & cost. Gaming devices were not of sufficient quality at that time and pin-arrays were unsatisfactory so the choice was reduced to research-grade 3d force-feedback devices. Only the cheapest such device, the crude Impulse Engine 3000, could be bought initially. Later, once the results had been produced from that, increased funding & decreased prices allowed the purchase of a higher quality device, a Phantom 1.0.

7.2. INTERFACING SOFTWARE TO HARDWARE

The software drivers for the IE 3000 & Phantom differed greatly. The former were simply IBM PC i/o bus addresses to read & write bytes to & from in DOS; the latter was system called Ghost which was a full 3d modelling language with Windows NT 4 drivers. These required totally different approaches.

The example software that came with the IE 3000 was inadequate due to faults in the encoder value roll-over handling, the conversion of encoder values to Cartesian position & hogging of processing time so that even the keyboard input was suspended. So little was useable that the driver was totally rewritten and the correct co-ordination (see Appendix A) imposed. To allow other processing to be carried out simultaneously, DOS was upgraded to Windows 95 (Windows NT does not allow direct i/o bus access) & the force-feedback loop written as a independent thread. The bespoke algorithms described below were created for object & texture simulation.

The Phantom's Ghost library was the other extreme. It was too comprehensive for the raw control needed so, instead of specifying objects to simulate, only a general-purpose force-field was specified and its call-back function used to divert control to the bespoke algorithm.

7.3. SIMULATING SOLIDS

There were several reasons for not using the comprehensive Ghost library:

- Ghost, although excellent for solids, does not readily support good texture simulation.
- The same system had to work on multiple hardware platforms, not just Phantoms.
- Ghost was not available at the start of the work because the IE 3000 was purchased first.

Instead a straightforward algorithm was developed. Its requirements included:

- Ability to simulate textured solids.
- Calculation time substantially less than 1 ms.
- Stability & robustness.
- Physical simplicity & a readily understandable basis to deter claims that our results were due to algorithm peculiarities.
- Ability to work on the mechanically poor IE 3000 not just the low friction low inertia Phantom.

The algorithm devised was based purely on a forcefields. Forcefields are solely functions of position not time, speed or path. This avoided complications from irregularities in the Windows clock and facilitated quick calculations. Touching was modelled as a point contact. An extended blunt contact might give a more intuitive feel but would involve performing far more calculations on each cycle of the simulation loop. More importantly, it would also reduce the finesse with which it would be possible for users to examine finely textured objects and add a complication to interpreting the psychological results.

The algorithm started with a smooth surface. Even simulating this was not trivial with problems like stability & managing with a limited maximum force.

To simulate a smooth textureless planar surface of a solid, there must be a normal force towards the surface when the handpiece is in a position corresponding to being inside the solid. When outside the solid, this force is of course zero. The spatial transition between the non-zero force inside and the zero force outside must be gradual to prevent the handpiece being oscillated in and out of the surface when the user tries to gently touch the surface.

A sensible dependence of the reaction force, R , on the displacement of the handpiece from the surface had to be chosen. The obvious simple realistic force profile would be to have it obey Hooke's Law, with the reaction force proportional to depth of the contact point below the surface, but the limited forces available from the

hardware made this infeasible for all but the softest objects. Instead, the surfaces were given a simulated elastic skin layer in which Hooke's Law is obeyed and beyond which the reaction force remains at a constant value. This was mechanically equivalent to not having a plane with a simple spring beneath it (Figure 15a) but instead a weighted lever linked to the plane by the spring (Figure 15b). After a certain amount of compression of the spring, the weight lifts from the floor and the force no longer increases. This gave a convincing haptic impression of a smooth surface. The thickness of the elastic skin or the force limit can be varied to change the object's apparent hardness. However, a large discontinuity in the force-field in a simulation that has discrete time-steps and little damping can lead to instabilities; this sets a limit to the minimum acceptable skin thickness, which in our case is approximately 0.1 mm.




Figure 15: Mechanical equivalent of hard surface model (a) without & (b) with force limit.

The force-field method of simulating a flat smooth plane is readily extensible to other primitive solid shapes which can then be combined to form haptic scenes. For each shape, a force-field was designed so that, at any point within the shape, the force acted outwards towards the nearest point on the shape's surface. For a cube, the force-field consisted of six square-based pyramids with the force in each acting towards the pyramid's base (Figure 16a) and for a sphere the forces were radially outwards (Figure 16b).




Figure 16: (a) Cube & (b) sphere force fields.

The different primitive haptic shapes (and textures) naturally fitted within an object-orientated scheme for the computer programming. See Appendix B for further information.

All the haptic objects could be rotated and translated by specifying a rotation matrix, \underline{M} , and a translation vector, \underline{v} . This was allowed for by replacing the force, $\underline{F}(\underline{r})$, calculated for a position \underline{r} with \underline{F}' defined by

$$\underline{F}'(\underline{r}) = \underline{M}\underline{F}(\underline{M}^{-1}(\underline{r} - \underline{v})). \quad \{1\}$$

This worked by untranslating then unrotating the position vector to a standard position before using $\underline{F}(\underline{r})$ to calculate the force then rotating the force back to correct orientation. Since all of this, except the function \underline{F} , uses the same calculation for all shapes, the algorithms that calculate the forces for each different haptic object do not need to take rotations and translations into account independently.

Matrix transformations include reflections, shearing and scaling (which may be different along the three axes) as well as rotations. To allow it to work with any non-singular matrix \underline{M} , not just the orthogonal ones that represent rotations and reflections, the covariant nature of the force vector compared to the contravariant nature of the position vector had to be taken into account. This required the matrix that was used to rotate the force to be the transpose of the inverse of the matrix that was used to rotate the co-ordinates [Mathews & Walker 1970]. Equation 8 then becomes

$$\underline{F}'(\underline{r}) = \left(\frac{|\underline{F}(\underline{M}^{-1}(\underline{r} - \underline{v}))|}{|\underline{\tilde{M}}^{-1}\underline{F}(\underline{M}^{-1}(\underline{r} - \underline{v}))|} \right) \underline{\tilde{M}}^{-1}\underline{F}(\underline{M}^{-1}(\underline{r} - \underline{v})) \quad \{2\}$$

where the scalar prefactor expression is just to ensure that the magnitude of the force is not altered by the transformation.

Although creating a new class and algorithm for each new species of shape is acceptable for the primitive shapes like cuboids, flat planes, spheres and cones, it is not practical for more complicated shapes. However, more complicated shapes can

be constructed by combining instances of the primitive shapes. For example Figure 17 shows a simple model of an armchair composed of nine primitive shapes grouped together. The back, seat and arms are cuboids with the fabric represented by a sinusoid texture (with short-period low-amplitude sinusoids in two orthogonal directions multiplied together to produce fine bumps instead of ridges). The top surface of the seat cuboid is made softer than the other surfaces to represent the cushion. Non-orthogonal matrix transformations have been used to create the short legs by compressing spheres in the vertical direction and the footrest by shearing a cuboid.




Figure 17: Chair simulation composed from basic units.

A group of objects could itself be treated the same way as any other object including being rotated, translated and used as a single object in further groups of objects. For example the chair example above could be duplicated, the duplicates first rotated then moved to different positions before being combined in a scene with a haptic model of a coffee table. In this way complex scenes can be built up as a hierarchy of groups starting from very simple primitive shapes.

In order to save the program from having to query every object in a scene to find out which is being touched, an extra function is included in the algorithm for each haptic shape. This calculates a bounding sphere outside of which no part of the shape projects. A group can then quickly screen which of its constituent objects need to be examined when it is queried about the force at a point. Because the group will also use these bounding spheres to calculate its own bounding sphere for the use of groups of which it itself is a constituent, the force can be calculated from a well-arranged hierarchy of groups reasonably efficiently; time required increasing logarithmically with the number of basic objects in a scene. Scenes of ~100 objects in

a few levels of grouping have been found to update only 2 to 3 times slower than a scene of single object. Since the latter can update at approximately 14 kHz even when running as a background task (Windows 95, 200 MHz single Pentium-Pro PC, IE 3000) and force-feedback need only be updated at $\sim 10^3$ Hz [Minsky 1995], this is ample.

7.4. SIMULATING OF TEXTURES

The most convincing way of simulating textures as force-fields was found to be simply to represent a texture as a spatially varying local displacement of the surface out of its plane. Firstly the component of the reaction force that was normal to the plane, R , was calculated as if the surface were smooth except that in calculating the distance from the surface, the texture was taken into account by subtracting the local surface displacement from that distance. That is, for the case of an x-y plane with local surface displacement $s(x,y)$, by using $R(z-s(x,y))$ instead of $R(z)$. So that the total force, \underline{F} , was normal to the local texture of the surface (Figure 18) instead of the overall plane of the surface, force components are added parallel to the plane. In the x-direction this had to equal the normal component multiplied by $-\partial s/\partial x$ so that the gradient of the force in the x-z plane was $-1/(\partial s/\partial x)$ which was perpendicular to the local surface texture gradient, $\partial s/\partial x$. The y-direction component was similarly calculated. Consequently, the expression for the force was simply

$$\underline{F}(x, y, z) = -\frac{\partial s(x, y)}{\partial x} R(z - s(x, y)) \underline{\hat{x}} - \frac{\partial s(x, y)}{\partial y} R(z - s(x, y)) \underline{\hat{y}} + R(z - s(x, y)) \underline{\hat{z}}. \quad \{3\}$$




Figure 18: Mechanical equivalent of texture algorithm.

Fortunately, this form of force field is intrinsically conservative so the chances for feedback instabilities in the simulation are minimised. The texture was varied by specifying different displacement functions (e.g. sinusoidal ridges, grids of lumps, random surfaces etc.), amplitudes and periodicities. Unlike the textures simulated by the 'Sandpaper System' [Minsky 1995], which used two dimensions of force-feedback, these textures had forces normal to the plane as well as in it and allowed the user to move off the plane. The method of texture reproduction used here mathematically degenerates to that used in the 'Sandpaper System' in the two dimensional limit where $F(z)$ is ignored.

The magnitudes of the forces and the addition of textures to the shapes were performed in the same way as for flat planes though, of course, one must choose a suitable projection for mapping plane textures onto surfaces with non-Euclidean metrics such as spheres.

7.5. TIMING ISSUES

Converting an algorithm into a functioning program takes far more time, though less intelligence & originality, than designing the algorithm itself but leaves little worth reporting in a scientific report such as this. It is as separate job from both the physics and the psychology*. However, one aspect is worth reporting: how the vital timing was achieved when programming for the IE 3000 which, unlike the Phantom, did not come with a millisecond timer.

The first problem was to have the force feedback loop cycling sufficiently often without preventing the other aspects of the programs from running. The following methods were tested. The fourth was found to be satisfactory and was used:

- **Simple program loop:** worked but did not allow for keyboard input, GUI updates etc. in the same program.
- **Windows timer:** the official Microsoft way of performing a regular loop. It is a call-back function that can be triggered from the Windows message queue. Even on fast machines it is too slow with typical calls at < 25 Hz.
- **Idle-time processing:** fast but in time chunks with unpredictable & potentially long delays inbetween.
- **Multitasking:** a simple program loop that runs as a separate task and checks shared flags for instructions from the rest of the application.

* It was originally planned for the project to have a dedicated programmer but funding changes removed that post so the current author learnt C++ & the MFC Windows API then took on the rôle of programmer as well as physicist.

The second problem was that some experiments needed feeling times accurately measured but it was found that Windows clock, although of millisecond resolution, was only updated at irregular intervals averaging 15 ms but often several times longer. This problem was solved by using a loop counter as the timer and repeatedly calibrating it against the Windows clock when that was updated.

This technique was also used to control the frequency of the force feedback updates themselves. The updates need not have been regular, provided they were often enough, because the algorithm was time-independent but too fast rapid updates caused an unexpected problem. Once, when the computer was upgraded, the IE 3000 emitted a high pitched wail annoying to both experimenter & subjects. It was traced to updates being sent to the hardware at 14000 times per second. This was brought back down to 1000 per second by skipping a fraction of the updates. The fraction was determined from the calibrated loop counter and dynamically adjusted (with suitable damping) to maintain the desired average update rate. This was an analogue electronic engineering technique called a phase-locked-loop adapted to digital software.

7.6. TEXTURE EXPERIMENT APPLICATION

This application simulated textures on a horizontal plane which subjects would feel and subjectively assign roughness values to. The application included both a graphical editor for the textures and an automated system for running the experiments and logging the responses.

The appearance of the editor was based loosely that of Minsky's 'Sandpaper' system [Minsky 1995] where each texture is represented by a sub-window within the application window (Figure 19a). These could be moved, resized, duplicated, renamed, stored in files and, most importantly, edited. Editing was done via a dialogue box (Figure 19c) where the surface hardness, the spatial periods, texture amplitude etc. could be set. The texture patterns themselves could be ridges or lattices of sinusoidal, triangular or trapezoidal waves.

Any of the textures could be activated for haptic display simply by making its visual sub-window the active one. This method was chosen instead of allowing the handpiece to run over a simulation of all the texture areas as they were arranged on the screen so that the experimenter, not the subject, controlled what was being displayed. Moreover, for experiments, the visual clues could be removed by blanking out the drawn textures (Figure 19b).

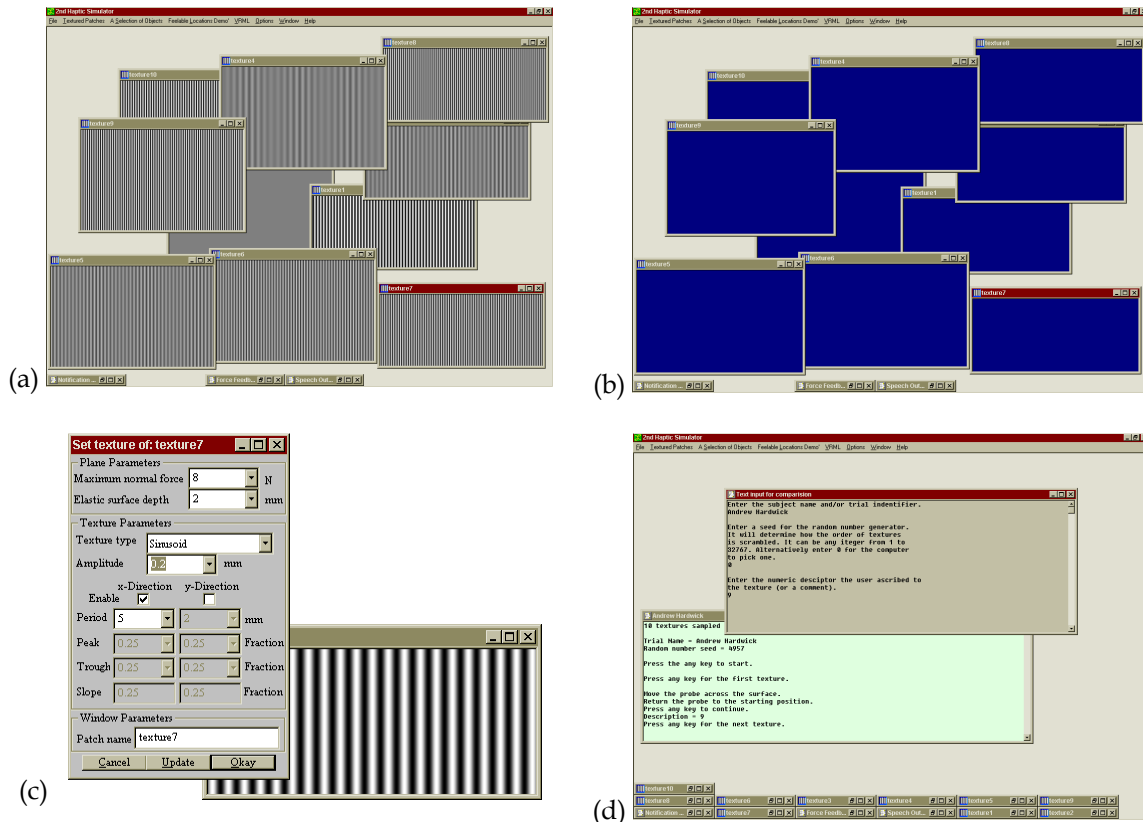


Figure 19: Screen shots of the texture experiment application: (a) experimental textures displayed; (b) the same textures blanked; (c) texture setting dialogue box with changed texture; (d) experiment in progress with texture windows minimised. To show detail (c) is at double the scale of the other images.

For the main experiments, the presentation of textures was automated (Figure 19d). The specified textures would be presented in a random order & for each:

- A beep would tell the subject the texture was ready.
- The subject would scan haptically across the texture from left to right.
- Texture disappears when the right hand side was reached.
- The subject would be asked for a subjective roughness.
- The haptic device would be forced back to the left for the next texture.

Once all the textures in a set had been felt the subjective roughnesses would be saved to a text file (Figure 20) along with the times spent feeling and parameters to identify the textures.

BT HAPTIC SIMULATION : TEXTURE COMPARISON TRIAL REPORT				
Trial title =		s17d		
Number of textures =		10		
Random number seed =		6		
Count	Description	Feeling time / ms	Texture number	Texture name
0	8	9119	5	texture5
1	8	2870	3	texture3
2	8	2250	9	texture9
3	16	2566	6	texture6
4	16	2992	7	texture7
5	4	2951	0	texture1
6	2	3034	8	texture8
7	1	2497	2	texture2
8	.5	3033	1	texture10
9	.25	3034	4	texture4

Figure 20: Typical output file.

7.7. SOLID OBJECTS EXPERIMENT APPLICATION

Compared to the texture experiment application, the solid objects one was simple with no visual representation of the objects was on screen and no automated data recording. The experimenter sets up a selection of objects from a dialogue box (Figure 21). Options included: sphere or cube; inside or outside presentation; size; surface hardness; rotation; & shear. Once set up, the simulated objects were switchable with single keypresses allowing the experimenter full control over the presentation order without giving significant clues as to the nature of the object to the subjects.

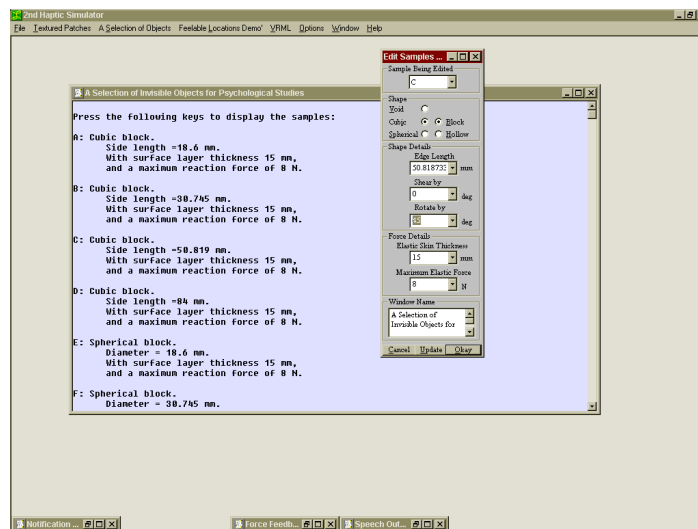


Figure 21: Screen shot of the solid object experiment application showing a settings dialogue box and a listing of set-up objects.

The ability to choose inside or outside presentation was a fortuitous accident. Originally created as temporary measure to work around a fault in the shearing mathematics, it was left in as an extra option &, became as factor tested in the

experiments resulting in a significant discovery (to be revealed later in this dissertation).

7.8. DEMO' APPLICATIONS

Demonstration applications were needed not just for publicity and to satisfy public interest but for casual opportunity-testing.

Both the above systems - textures and solids - also functioned as demonstrations. The barely discernible textures from the experiments were replaced with readily discernible textures and they were displayed graphically on screen (Figure 22).

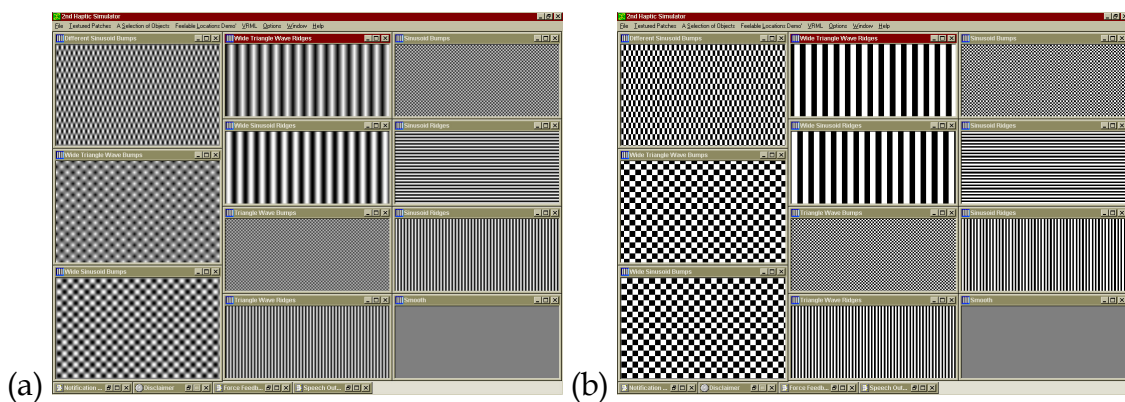


Figure 22: Screen shot of demo textures in (a) normal greyscale mode & (b) high contrast black-&-white mode.

Of course, BT advertising was also added (Figure 23).

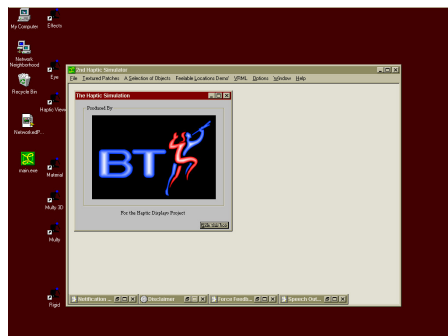


Figure 23: Screen shot of a BT advert in demo' application.

One extra application was created, consisting of over 40 haptically simulated scenes. These were independent but the user could move between them in the north, south, east, west as in traditional text-based computer adventure games using the keyboard. This format was chosen because it avoided the complication of using the same haptic interface for feeling objects and moving between locations. When

moving into a room (i.e. changing displayed objects), a short description was displayed as text on the screen and output verbally from a speech synthesiser. The rooms were arranged in three groups: "The School" which introduced basic shapes, positions, textures etc.; "The Modern Art Gallery" which used geometric exhibits to allow the user to practice feeling more involved scenes of multiple objects, textures & hardnesses; and "The Furniture Shop" where complex haptic representations of real (but scaled down, of course) objects of furniture were feelable.

The demonstrations that came free with the Phantom were also shown but those that came with the IE 3000 were not of useable quality.

8. RESULTS OF USER TESTS

8.1. INFORMAL EXPERIMENTS

Casual observation of people trying the haptic system (Figure 24) and listening to their comments - including both praise and complaints - lead to many interesting discoveries.



Figure 24: A visitor to a BT demo' using a Phantom.

The *ad hoc* nature of these discoveries unfortunately does make the following read like rather like a disordered list:

8.1.1. PRIOR HAPTIC EXPERIENCE HELPS

People who were used to working by touch were much quicker than average at getting used to recognise the simulated objects. They appeared to have more flexible and refined haptic exploration skills. This was most dramatically shown by the occupations of the 3 people who have been able to do so immediately: teacher of Braille; blind piano tuner; and glass carver. It was not simply linked to visual ability (the glass carver was sighted) but to experience with working by touch (glass is carved underwater to inhibit cracking).

The existence of such people who have virtually no problem with using the systems is also a reassurance that our simulations do work technically and therefore we really were examining human perception not merely inadequacies in our experimental system.

It is also shown to a lesser extent in that engineers used as casual test subjects during development at the BT Labs appeared to get used to the system quicker than the psychologists used as sighted subjects in the SDRU trials. This might be because the good spatial imagination which required for their jobs is also extremely helpful in building up a mental picture of objects felt one point at time. Alternatively it might simply have been that the engineers had more realistic expectations of the hardware.

8.1.2. PERCEPTUAL CONTEXT HELPS

When subjects were asked to identify complex models like the items in the furniture store from haptic simulations, most could not do so although they could detect the components of them. However, when prompted with clues that put the model in context such as “Can you feel an item of furniture there?”, many more did so. This is a common effect that has even been found in musical GUI outputs [Alty & Rigas 1998].

8.1.3. DIFFERING MENTAL MODELS

A haptic simulation can be visualised in different ways by different users. When users were asked to touch the ‘front’, ‘top’, ‘back’ etc. faces of a cube simulated in the IE 3000, it was unexpectedly found that different people touched different sides. Further questioning revealed that some imagined the cube to be at their finger tips where the finger touched the rod (Figure 25a) but others imagined that the rod was a lever that pivoted about and slid through the mechanism’s midpoint and it was the other end of the rod that touched the cube (Figure 25b). Later it was found that even more alternatives are possible because a user’s mental model may not be a physically realisable one but instead be only fragmentary [Hammond *et al* 1983]. For example the user might refer to left-right motions and positions as if the simulated object is in the space outside the mechanism yet refer to up-down ones as if it were inside.




Figure 25: Alternative mental models for the haptic environment: (a) finger directly touching object; (b) handpiece as tool touching object.

From a psychological perspective, it was something to investigate. Of 19 subjects from Solid Object Experiment 1 below, 14 imagined the objects inside, 4 outside, and 1 mixed. 3 of the 4 who imagined it outside were blind suggesting that the mechanical appearance of the hardware encouraged the interior mental model.

From an engineering perspective, it was a problem needing a solution. For the IE 3000 experiments it, the solution was to adapt the kinematic equations (Appendix A) to represent the object being inside the mechanism and allow the experimenter to choose between these or the originals after checking where the subject naturally considered 'up' & 'left'. Of course, this still left the ambiguity of how long subjects imagined the probing stick to be & did not help the few with physically impossible mental models. Fortunately the Phantom mechanics presented much less ambiguity.

8.1.4. IMPOSSIBLE REACHING

It is possible to reach through objects whilst touching them (Figure 26). This unnatural action cannot be prevented because the system can only measure the position of the point contact not the where the rest of the user's hand & arm are. It was initially a concern that this would detract from the simulation but, in practice, people generally had no qualms with reaching through objects to feel object backs. Instead of being considered an unnatural drawback, it was considered an advantage

because it enabled feeling of parts of models that would be inaccessible in reality. We named this effect 'Impossible Reaching'*.




Figure 26: Feeling the same point in a haptic simulation via routes that are (a) possible & (b) in reality.

Note that not everyone spontaneously felt behind objects. This was not so much because of abhorrence of Impossible Reaching but because they were so used to computer displays being 2 dimensional that they needed prompting before they would intentionally move the device in & out at all. One they were using it 3 dimensionally, Impossible Reaching was again generally accepted without comment.

8.1.5. OBSERVATIONS AT DORTING COLLEGE

Both the IE 3000 & Phantom systems were demonstrated at the Access '98 day at Dorting College which is a Sixth Form level college in Seal near Sevenoaks in Kent run by the RLSB. Concerns (by the current author) that the experimental systems were inadequate for public demonstration proved unfounded because the demo's were very well received by the many blind (& most of the sighted) visitors to the stand. The only exception was one child, with a tactile disability as well as blindness, who could not tolerate vibrations.

The show was also valuable for the information, ideas & encouragement received from the intelligent, imaginative & enthusiastic blind people met (Figure 27). The main negative comment was not about the potential utility of such systems,

* The name 'Impossible Reaching' was originally a crude punning quip but name stuck & was fixed in print.

the systems on show or the fact that it was not yet a finished product but just the valid complaint that high quality haptic outputs were far too expensive.



Figure 27: Blind visitors to our stall at Dorting College (a) using the Phantom, (b) using the IE 3000, (c) being told about the project by Stephen Furner whilst reading a Braille show guide & (d) amused by discovering “a dump in the middle of the floor!” in the a Phantom demo’.

8.1.6. LEARNING

As expected, people show learning effects. There is often a *aha!* [McCrone 1993] effect when a haptic object is first comprehended. More subtly, people tracing the

contours of solid objects initially overshoot the but in a few minutes the paths settle down to tracing in contact around the objects smoothly.

8.1.7. 'NO PROBLEM'

There were several aspects to the simulations that were thought, in advance, to be psychophysical inadequacies but were not found to cause much of a problem:

- Stimulation of the wrong haptic receptors: Naturally when feeling a texture it is by tactile nerves that sense skin deformation & when a mechanical apparatus moves a limb it is proprioceptive sensing of joint positioning that is used. Our systems use mechanical apparatus to display textures so it is combination of the wrong haptic receptors for texture being stimulated and correct receptors being stimulated only by the indirect effects of the vibration and pressure of the handpiece on the skin. This is complicated but, fortunately, people have no particular difficulty feeling textures this way. Maybe this is because it similar to tool use.
- Orientational transposition: There is a complication in that computer users are familiar with a mouse moving on horizontal table. This they consider to match a pointer moving on a vertical screen. With a 3d system it is possible to remove this transposition; so should a 3d system imitate the 90° rotation of 2d systems or not? In practice, users familiar with computers often defaulted to transposition unless the image onscreen clearly imitated 3d but people could easily switch between modes when asked.
- Limited force: Even the forces limited to ~8 N from the IE 3000 were found adequate, with suitable force-feedback algorithms, to give the impression of hard surfaces*.

8.1.8. HIGHER ACCURACY NEEDED FOR HAPTIC THAN VISUAL PICTURES

Much haptic simulation research is into making simulations more realistic because users complaint if the feeling is not perfect. This is in contrast to visual simulations where people accept jerky, low resolution & incorrectly shaded 'VR' simulations composed of very simplified badly joined 3d blocks with implausible mechanical responsivity displayed on 2d screens. They even accept cartoons, impressionist paintings, etc.!

Maybe this is because vision is better at providing a broad overview so irregularities can be disregarded or compensated for. Alternatively, it might just be

* Except for the few people who deliberately set out to find the hardware limits or break the device (thankfully the Phantom includes an overheating sensor).

that people are used to experiencing visual representations from cave paintings to TVs whereas haptic simulations are an almost totally new experience.

8.2. NOTES CONCERNING THE STATISTICAL ANALYSES

The data from the major studies (in terms of time taken in data collection) were analysed by an ANOVA program. ANOVA is an automated system that uses Fischer's Analysis of Variance to detect dependence of variables upon parameters from a data set. It essentially works by: grouping data by parameter values varying only one parameter at a time; calculating the variance within those groups; comparing that variance with the variance across the groups using the 'F' test to calculate significance based on the degrees of freedom involved; combining groups that are not significantly different; repeating if necessary; and reporting the relationships found and what their significance is. It can then go on to find non-linear dependencies in the data from the combined action of two or more parameters by subtracting predictions made from the previously found relationships from the data and reanalysing the residual against parameter value combinations. It is frequently employed by psychologists because it is a reasonable general purpose analysis method that saves them from having to understand enough mathematics to devise analyses optimised to particular experiments. It is also robust in the sense that its naïve application of ANOVA is unlikely to give false positives yet can still detect dependencies, just with less sensitivity than an optimised analysis. For example the length parameters used here were expected to have approximately proportional, or at least monotonic, effect on the subjective responses but that *a priori* information was not used because ANOVA simply treats the parameter values as labels (there is a variant called ANCOVA, based on Analysis of Covariance, that does use this information). However ANOVA still detected the relationships reported below.

The data from the major experiments were analysed in a formal 'production line' manner (computational method chosen in advance and applied scrupulously) rather than a dynamic 'exploratory' manner (intelligently adapting the statistical techniques in light of results appearing). This is standard practice in psychophysical research because, although it is likely to miss any discoveries not specifically searched for & is inefficient in the number of trials needed, it avoids the need to distinguish between *a posteriori* & *a priori* hypotheses in the statistics.

The data analysis in those experiments was also left to the end of all the trials and interim results were not obtained to feed back into optimising experimental methods during the course of the trials even though as simple a procedure as stopping trials once the required confidence level has been reached can typically

halve the number of trials needed [Brigham 1989]. There were two main reasons for continuing. Firstly, changing the parameters during the study unbalances the design making the statistical analysis by ANOVA less easy. Secondly, psychology experiments – even innocuous ones like these – have to be ethically approved in advance which limits the scope for adaptation after ethical sign-off.

The number of subjects was typically chosen to be greater than 20 even when the data from the first few subjects showed clear results with high significance. The extra subjects were still needed to show that the results apply to the general population. The conventional use of $p < 0.05$ for generality implied over 20 different people had to be tested. Of course, psychophysical phenomena tend to be much more uniform across different people than, for example, social psychology phenomena so one has a high *a priori* probability for homogeneity which could be used to reduce the number of subjects needed. Bayes' theorem enables the mathematical incorporation of such *a priori* probabilities complication but the use of less than 20 subjects is still, regrettably, something which hinders getting one's results published in mainstream psychology journals. Academic psychology departments require refereed publications for funding and therefore are usually obliged to waste time repeating experiments on unnecessary numbers of subjects.

Here is not the place to go further into the merits of 'production line' verses 'exploratory' data analysis, consistency verses flexibility in experimentation or journal requirements but it does explain why the some of the results below are quoted with chance probabilities so low that, with hindsight, the reader may wonder if the experiments could have been performed with fewer trials freeing time for additional investigations. They could have been but the option was not easily available to academic psychologists for other reasons.

8.3. SOLID OBJECT EXPERIMENTS

8.3.1. SOLID OBJECT EXPERIMENT 1

The whole of the object size data set was analysed together by ANOVA. The results are in Table 1. The two missing cube values were omitted from the experiment for practical reasons: the smaller one was too difficult to locate in space & the larger clipped the IE 3000's workspace. Although the actual values of the perceived sizes are distorted by the calibration error (which saturated results at the larger end of the scale because the simulated objects were larger than the multiple choice pictures), it was clear that objects were larger when felt from inside than from outside. This haptic illusion was an unexpected & counterintuitive discovery which

we named the ‘Tardis Effect’ after the TARDIS time machine / space ship in the popular BBC series ‘Dr. Who’ [Parkin 1996] which was larger inside than outside.

Shape	Actual Width /mm	Perceived External Width / mm	Perceived Internal Width / mm
Cube	18	(not used)	18 ± 4
	25	16 ± 5	17 ± 3
	36	20 ± 5	24 ± 2
	45	24 ± 7	(not used)
Sphere	25	12 ± 4	21 ± 1
	36	18 ± 5	23 ± 1
	45	23 ± 8	25 ± 1

Table 1: Mean perceived sizes of objects.

The rotation and corner angle results are in presented in Table 2 & Table 3 but they did not show as interesting a discovery, merely that the subjects were generally not accurate at recognising angles and almost hopeless with rotation angles in spite of it being multiple choice and there being a clear “top slopes down left” verses “top slopes down right” distinction between the 30° & 70° cases. The rotation difficulty may have been because the objects were free floating in space giving little in the way of a reference.

Actual Rotation	Perceived Rotation
30°	40° ± 12°
50°	52° ± 12°
70°	48° ± 18°

Table 2: Mean perceived rotation of cubes.

Actual Angle	Perceived Angle
18°	20° ± 11°
41°	37° ± 11°

Table 3: Mean perceived corner angle of sheared cubic hollows.

No significant difference was found in this experiment between the blind and sighted subjects.

8.3.2. SOLID OBJECT EXPERIMENT 2

The small number of subjects & tests in this short experiment enables the whole data, not just the averages, to be presented in Table 4. Even with this small sample,

simplistic treatment & no statistical analysis, the Tardis effect is blatantly clear. In all but one of the 12 cases, the internal widths of the cubes was felt to be greater than the internal width.

	Subject 1			Subject 2			Subject 3			Subject 4		
Actual	v	he	hi	v	he	hi	v	he	hi	v	he	hi
19	15	10	20	15	31	33	20	13	20	25	13	25
31	30	20	100	25	33	57	30	25	45	25	25	38
51	50	50	150	45	58	64	55	100	60	38	51	76

Table 4: Perceived widths of 3 cubes by 4 subjects visually (v), haptically externally (he) & haptically internally (hi). All lengths in mm.

The visual estimates of sizes confirmed that subjects 1 to 3, who were real engineers, could estimate lengths reasonably as required for this experiment. However, even subject 4, a publicist only nominally an engineer, showed the Tardis Effect clearly despite giving poor visual estimates and using a mixture of cm & inches yet, when asked, getting the number of mm in an inch wrong. The Tardis Effect is indeed such a strong & reliable haptic illusion that it could be demonstrated as a ‘party trick’ like many famous visual illusions frequently are.

8.3.3. SOLID OBJECT EXPERIMENT 3

The object size data set was once again analysed together by ANOVA. The results equivalent to Experiment 1’s Table 1 are shown in Table 5 & Figure 28. The Tardis effect is clearly apparent ($p < 0.001$). With the calibration problem solved, the relative discrepancy between perceived and real size becomes interesting. For external presentation, the relative discrepancy increases with increasing size but for internal presentation it is constant. The relative discrepancy is remarkably similar for the two different shapes ($p = 0.67$ for difference being by chance), especially if averaged over sizes as in Table 6.

Shape	Actual Width /mm	Perceived External Width / mm	Relative Discrepancy	Perceived Internal Width / mm	Relative Discrepancy
Cube	20	12 ± 8	-38 %	20 ± 10	-2 %
	30	17 ± 7	-43 %	25 ± 10	-18 %
	40	21 ± 9	-47 %	29 ± 12	-28 %
	50	27 ± 12	-46 %	35 ± 14	-30 %
Sphere	20	11 ± 6	-45 %	18 ± 9	-9 %
	30	19 ± 9	-37 %	25 ± 11	-15 %
	40	24 ± 14	-39 %	30 ± 15	-24 %
	50	23 ± 12	-54 %	37 ± 16	-27 %

Table 5: Mean perceived sizes of objects & relative discrepancy.




Figure 28: Mean perceived width (averaged over cubes & spheres) compared to actual width.

Shape	External	Internal
Cube	-44 %	-20 %
Sphere	-44 %	-19 %

Table 6: Mean relative size discrepancy.

The varying of object hardness was to test the hypothesis that the Tardis Effect was due to the user pushing into the simulated objects because of the limited force available. This is equivalent to a real object being bent inwards when felt from the outside and bent outwards when felt from the inside (Figure 29). If this hypothesis was correct, the Tardis effect should be stronger for softer objects. However although varying the surface reaction limit did affect subject's subjective rating of the hardness very clearly ($p < 0.001$) & almost linearly, it had no significant effect ($p = 0.29$) on the perceived size.




Figure 29: A disproved hypothesis for the Tardis effect: surface deforming (a) inwards from outside but (b) outwards from inside.

Once again, no significant difference ($p=0.76$ for the observed difference being from chance) was found between blind and sighted subjects.

8.3.4. SOLID OBJECT EXPERIMENT 4

The object size data were analysed, as normal, by ANOVA. The results equivalent to Experiment 1's Table 1 are shown in Table 7 for both the stylus & thimble versions of the hardware. As has come to be expected, the Tardis Effect was unambiguously visible (& a, hardly necessary, F-test confirmed $p<0.001$). The main difference from the IE 3000 results of the previous experiment was that, rather than the effect being the same for cubes & spheres, spheres were felt to be larger than cubes of the same width. This effect is in addition to the Tardis Effect.

Shape	Actual Width /mm	Perceived External Width / mm		Perceived Internal Width / mm	
		Stylus Handpiece	Thimble Handpiece	Stylus Handpiece	Thimble Handpiece
Cube	27	19 ± 10	20 ± 10	34 ± 13	34 ± 9
	36	28 ± 13	27 ± 9	45 ± 16	45 ± 16
	45	37 ± 19	37 ± 14	52 ± 15	53 ± 18
Sphere	27	17 ± 7	19 ± 10	31 ± 13	27 ± 10
	30	22 ± 6	23 ± 7	40 ± 19	36 ± 11
	50	32 ± 12	28 ± 11	46 ± 17	45 ± 18

Table 7: Mean perceived sizes of objects felt internally on a Phantom with different handpieces.

The results of the corner angle investigation, Table 8, were not very interesting.

Actual Angle	Perceived Angle	
	Stylus	Thimble
18°	22° ± 11°	21° ± 10°
41°	36° ± 9°	33° ± 12°
65°	51° ± 10°	46° ± 15°

Table 8: Mean perceived corner angle of sheared cubes.

Yet again, there was no significant difference was found between blind and sighted subjects.

8.4. TEXTURE EXPERIMENTS

The data were analysed first within subjects & conditions to extract the Stevens' exponents (by linear regression fits of the logarithm of the perceived roughness to the logarithm of groove width). These were then processed through ANOVA as normal to detect relationships & calculate significance values.

8.4.1. TEXTURE EXPERIMENT 1

Of the 22 subjects, 14 showed significant ($p \leq 0.05$) relationships between perceived roughness and groove width. However only 3 of those 14 had the expected positive exponents and so had felt wider grooves to be rougher. The other 11 had negative exponents and so had felt the narrowest grooves to be roughest.

The division between blind and sighted was interesting. Every one of the blind subjects showed a significant relationship whereas only 5/13 of the sighted subjects did. All 3 of those subjects who showed the conventional negative exponent were blind.

A casual observation from this experiment was that, as in many psychophysical experiments, subjects were actually better at discriminating stimuli than they consciously thought they were. Many subjects who were frustrated at what they thought was their inability to detect difference in textures were found, once the data was averaged, to have a significant correlation of their guessed perceived roughnesses to the actual simulated groove widths.

For more detail see Colwell 1998b.

8.4.2. TEXTURE EXPERIMENT 2

Of the 23 subjects, 22 showed significant relationships between perceived roughness and groove width with at least one of the two handpieces and 18 showed

it with both. The result for the Phantom is substantially higher than with the less well engineered IE 3000. This time only 1 of those 22 showed the traditional positive exponent. Virtually everyone* felt narrower grooves to be rougher. Individual results from linear regression are in Table 9.

	Stylus		Thimble	
	Exponent	Significance	Exponent	Significance
Sighted per Subject	-0.71	$p<0.0005$	-0.83	$p<0.005$
	-1.1	$p<0.0005$	-1.6	$p<0.0001$
	-0.78	$p<0.005$	-0.47	$p<0.01$
	-1.1	$p<0.0005$	-1.3	$p<0.0005$
	-0.59	$p<0.01$	-0.89	$p<0.0001$
	-0.58	not sig.	-1.5	$p<0.0001$
	-0.69	$p<0.005$	-0.60	$p<0.0001$
	-0.70	$p<0.005$	-0.23	not sig.
	-0.65	$p<0.01$	-1.1	$p<0.0005$
	-0.35	not sig.	-0.56	$p<0.01$
	-0.39	$p<0.05$	-0.38	$p<0.05$
	-0.036	not sig.	0.10	not sig.
	-0.23	$p<0.0001$	-0.59	$p<0.0001$
Sighted Combined	-0.60		-0.77	
Blind per Subject	-0.83	$p<0.001$	-1.5	$p<0.0001$
	-0.42	$p<0.0005$	-0.60	$p<0.0001$
	-0.45	$p<0.01$	-0.70	$p<0.0001$
	-0.23	$p<0.05$	-1.5	$p<0.0001$
	-0.73	$p<0.0005$	-0.85	$p<0.0005$
	-0.40	$p<0.001$	-0.16	not sig.
	-0.52	$p<0.01$	-0.60	$p<0.0005$
	0.46	$p<0.0001$	0.61	$p<0.0001$
	-0.51	$p<0.005$	-0.77	$p<0.0001$
-0.43	$p<0.005$	-0.89	$p<0.0001$	
Blind Combined	-0.40		-0.69	

Table 9: Stevens' exponents (& significance levels) for perceived roughness as a function of groove width for individual sighted & blind subjects for 2 different Phantom handpieces. Also shown are the exponents from combined group data.

There was less obvious distinction between blind and sighted results because there were far fewer cases of no correlation or positive exponent found but the 2 cases found fitted with the previous results in that the no-correlation result was sighted & the positive-exponent one was blind. Formal analysis by ANOVA on the

* One could write "all subjects" instead of "virtually all subjects" if one cynically applies the simplistic significance limit common in psychology of anything greater than 19 out of 20 is certainty!

individual exponents from Table 9 revealed no statistically significant difference between the blind and sighted results but one between the stylus & thimble handpieces (Figure 30).

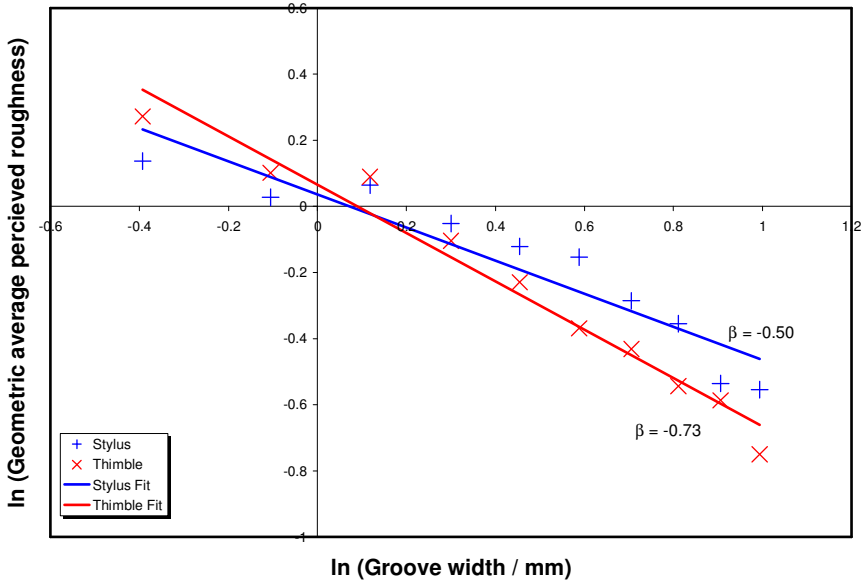


Figure 30: A log-log plot of roughness verses groove width for the 2 handpieces with linear regression fits.

For more detail see Penn 2000.

9. DISCUSSION OF RESULTS

9.1. GENERAL OBSERVATIONS

The discoveries listed immediately below are worth noting, and are important, but are not involved enough to justify having a whole subsection discussing them. Most of these may seem obvious with hindsight:

- **Physiological realism is not vital:** Impossible reaching, infinitesimal contact points & tactile receptors stimulated by mechanical motion instead of shape distortion of the skin were all acceptable.
- **More literal representation is needed for haptic than visual pictures** because most people are unfamiliar with haptic pictures and so compare them to reality not simplified pictures.
- **Alternative mental interpretations are possible for the same haptic environment.** The inside/outside alternatives for the IE 3000 are our most striking example. People did not consider that there could be an alternative to the version they automatically picked but most could switch at will when informed.
- **Blind & sighted people have the same haptic perception.** But blind people (& others used to working with objects by feel) get used to recognising simulated objects far more rapidly than & cope with low quality haptic systems better than others do.
- **A good spatial memory helps.** This might be an artefact of the point feeling mode but there is prior evidence that people synthesise shape information from separate touch measurements rather than all at once suggests this is general.
- **Some people are haptically disabled.** Not everyone can use a haptic output. Some are physically disabled (including a senior BT manager with severed nerve to his right forefinger) & some who find haptic mental visualisation extremely difficult (including one unlucky experimenter on this project for whom the experiments were very frustrating). Interfaces therefore must not be purely haptic any more than purely visual.

9.2. TARDIS EFFECT

The Tardis Effect is strong, easily repeatable & intriguing so it is rather surprising that it has not already been reported in the literature. A few other haptic size & shape illusions have been reported though. Many studies [summarised in Boff & Lincoln 1988] have shown a curvature illusion. When a finger traces along a curved path, the curvature is estimated as more convex with respect to the body than it really is. This curvature illusion has, unlike the Tardis effect, an obvious cause - the joints of a human arm naturally make it move on a concave path. Confirmation came from the illusion increasing when the arms' radii of curvature were decreased by

constraining the subjects' elbows. Another illusion is that people feel external angles to be more acute than they really are [Lakatos & Marks 1998] with the underestimation proportional to the angle size. Although angle underestimation could cause a Tardis effect for spheres if the subject estimated diameter from surface curvature, it cannot be the main cause because the Tardis effect also works with cubes where the angles are not related to size. A third illusion is that paths feel longer when radial to the body's axis than when transverse but that too cannot explain the Tardis effect because the same feeling directions were available for both internal & external feeling.

We have generated several hypotheses for the mechanism behind the Tardis effect including:

- **Elastic distortion:** Objects distort when pushed (in the force-field simulation algorithm, the contact point moves beneath the surface which has the same effect). Feeling from inside pushes them outwards increasing the size and from outside reduces the size. Solid Object Experiment 3 disproved this with different hardness objects.
- **Touching with fingertip centre:** The simulation places the surface at the finger centre. Naturally one feels with the finger surface & therefore subconsciously adds half a finger width. Solid Object Experiment 3 disproved this by showing the effect also worked when feeling by stylus.
- **Relative free volume:** When the subject is feeling the outside, they have a larger volume of free space to compare the object to than when they are feeling the inside so the object feels relatively smaller.
- **Maintaining contact outside:** When feeling outside an object, the user has to work to remain in contact with it and thereby traces a shorter path than when inside.

The 'relative free volume explanation' is from Paul Penn who will be testing it but the 'maintaining contact outside' one is from the current author so some more detail is presented here.

To avoid getting lost when feeling outside an object, people typically maintain contact with it which necessitates following a conservative path. The details are different for cubes & spheres but in both cases the interior path can be longer. For a cube, feeling edge to edge outside in contact with the surface gives a path length between the side length & $\sqrt{2}$ times it taking a face diagonal (Figure 31a) whereas a spatial diagonal is available from inside raising the maximum length to $\sqrt{3}$ (Figure 31b).




Figure 31: Possible paths traced when a cube is felt from (a) outside & (b) inside.

With a sphere, it is difficult to maintain an equatorial path outside so a smaller circle is typically taken (Figure 32a) whereas it is easy to maintain from inside simply by pushing out (Figure 32b).




Figure 32: Typical paths traced when a sphere is felt from (a) outside & (b) inside.

An experiment to test this hypothesis would be to record the actual paths taken as people feel objects & compare them to their size estimates.

9.3. THE PHYSICAL CAUSE OF ROUGHNESS

The negative Stevens' Law exponents found for roughness as a function of groove width in the texture experiments above require explanation because virtually all prior groove-width experiments reported in academic literature showed positive Stevens' exponents.

This suggested that the groove width itself was not the physical parameter which should have been used in fitting to Stevens' Law but instead some other parameter which is in turn determined by the groove width in way dependent on the experimental set-up. I.e. changing the groove width caused a concomitant change in another variable which then determined perceived roughness. The question raised is:

what is this other physical variable which really determines the roughness? Discovering that would be a very useful result for roughness simulation and a scientifically satisfying generalisation.

The use of groove width in these experiments was partly because it is intuitively obvious but was mainly because it had been shown to determine roughness in the real texture experiments of Lederman. Lederman in turn probably chose groove width to fit the earlier experiments of Stevens who used sandpaper grit sizes (an arbitrary scale approximately inversely proportional to grain diameter). However, the grain size of sand not only determines the width between bumps but also the bump height, stick-slip step forces, variability, etc.. Height is ruled out by our, & Lederman's, lack of height variation. Predictions from stick-slip forces, contrariwise, fit *every* study in sign of exponent. A further refinement is to normalise this parameter down to a friction coefficient because the studies which fixed the applied force [Lederman & Taylor 1972] showed that the magnitude of applied force was a minor factor (exponent of 0.13).

The proposed physical parameter determining the psychological feeling of roughness is essentially just "**The greater the jolts from passing over the bumps, relative to the pressure applied to the surface, the rougher it feels.**" in casual terms. More formally it is: "The dominate component of the haptically perceived roughness is proportional to the magnitude of the effective friction coefficient in the stick-slip motion across the surface raised to fixed power (the Stevens' exponent) dependent upon the physical set-up & the individual involved". Of course, other factors such as applied force do have an effect but those effects are less strong.

Here follows detail of how the hypothesis fitted the various studies:

- **Texture Experiment 2:** Only the period varied & simulated contact point was small enough to fall into each groove. Therefore the steeper sides of narrower grooves caused greater jolts so narrower felt rougher.
- **Texture Experiment 1:** Mainly as for Texture Experiment 2 but poorer quality hardware, especially in friction & inertia which reduced the ease of dropping into grooves, making feelings less distinct so many subjects could not discriminate textures at all. The few blind people who reported positive Stevens' exponents may have compensated for the difficulty of feeling textures spontaneously & instead built up a mental model of the surface shape & then reported wider as rougher. This are totally different psychological processes.
- **Lederman bare finger:** Fingers are wider than their grooves & spacings so they do not slip far down into spaces. Wider grooves allow further penetration therefore caused more resistance and bigger jolts therefore wider felt rougher.

- **Lederman-Klatzksy probe:** Wider was rougher, as with fingers, when probe contact diameters were greater than bump spacings. As spacings approached contact diameters, jolts no longer increased (the surface being flat between constant size bumps) and roughness was constant. The smallest probe showed a slight drop in roughness from the minor effect of simply less bumps being hit.
- **Minsky-Steele 2d simulation:** The width-based Stevens' exponent was not statistically significant [Minsky 1995 p. 98] because their grooves were almost rectangular not sinusoid so width did not effect jolt magnitude. The maximum magnitude of the horizontal force, which in 2d was equivalent to groove height, was the sole determinant of roughness [Minsky 1995 p. 105] with greater forces giving greater jolts and therefore feeling rougher.
- **Stevens' sandpaper:** Lower grit number sandpapers had bigger grains which caught the skin more giving jolts & therefore felt rougher.
- **Jansson-Green simulated sandpaper:** In Green's texture algorithm, lower grit numbers had greater vertical displacements giving greater reaction forces giving greater jolts for a given friction coefficient. Lower grit numbers also gave greater widths of the random friction distributions giving a greater upper. Both factors make lower grit numbers feel rougher.

9.4. HAPTIC GAMMA CORRECTION

Stevens' Law was found to hold well for the texture data per person but the actual exponent varied greatly between people. The psychological reason for this is still unknown but, for practical purposes, it is a something which might need to be compensated for.

The compensation could be achieved simply by raising the roughness determining parameter (see above) to a fixed per-person power because Stevens' Law is a power law. This is identical to the mathematics of 'gamma correction'* used in television sets & computer monitors so I provisionally called it 'haptic gamma correction'. Just as in the video case, the correction can be partly made in the recording & partly made in the playback. In television the recording correction

* The need for video gamma correction stems from physics; the brightness from the screen phosphor on the screen being proportional to the voltage applied to the electron gun at the back of the tube raised to a fixed power γ (gamma). Naturally $\gamma \approx 2.5$. It could be compensated for completely (to $\gamma = 1$) & concern the user as little as other electronic details if it were not for typical television use & cheap PC design. Televisions are typically used with dark surroundings which distort brightness perception which makes $\gamma \approx 1.4$ look more correct than $\gamma = 1$. Typically $\gamma \approx 1.8$ because the extra 0.4 makes colours look artificially bright which viewers like. [Aside: increasing γ on dreary 'Eastenders' makes it look like jolly 'Neighbours'!]. As for computers, original PCs had cheap screens so $\gamma \approx 2.4$; computers often used for graphics like Macintoshes & Silicon Graphics had better correction (default $\gamma \approx 1.8$ & 1.4 respectively). A legacy of that early cheap decision is that WWW pages have poor cross-platform colour fidelity.

dominates because it was cheaper to do it in one camera than in many television sets in the early days of electronics. In haptics, a sensible division would be any correction for general difference between the Stevens' exponent for the method of recording & reality to be made in recording bringing the exponent to some agreed standard value followed by a further correction in playback dependent on the simulation method and the individual user.

However, haptic gamma correction might be unnecessary because the same perceptual variation is likely to apply to real textures. It depends on whether one desires to simulate a texture emulating a particular real texture or whether one desires to so give a feeling of a particular roughness. Even if not necessary, haptic gamma correction could be an adjustable user preference akin to a contrast control which one could turn down to reduce strain or turn up when working in haptically noisy environments. It could even be used to augment tactile displays for haptically impaired people.

10. RECOMMENDATIONS

10.1. BT HAPTIC RECOMMENDATIONS

From the end of the Introduction until here, this paper has concentrated on communal academic results but business recommendations for BT can also be drawn from the work. The following are deliberately general & qualitative because BT's current use of computational haptics is so low.

- Ensure its networks are suitable for haptic transmission. Transmitting haptics will become a requirement from customers.
- Remember that haptic outputs need user customisable settings, e.g. haptic gamma correction & alternative mental models, like audio-visual ones have contrast & volume controls. Designers must not assume what feels best for themselves suits everyone.
- Make use of haptics in user interfaces where it increases ease of use or productivity. This applies both internal (e.g. to increase operator speed) & in products.
- Patent haptic gamma correction (this is underway).
- Carry out further haptic research into the fundamentals of the human sense of touch which vitally underlies any more specific investigation such as into network requirements.

10.2. BT OTHER RECOMMENDATIONS

Two recommendations from this work are not specifically haptic:

- Ensure BT services are accessible by disabled users. The prime example is its main public WWW site *[redacted business information]*.
- *[redacted business information]*.

10.3. FURTHER RESEARCH

Computational haptics is a growing field with much scope for experimentation. Even keeping solely to obvious experiments to check out hypotheses from the above work there are:

- Test for Tardis Effect with real objects (this is underway).
- Check effect of constrained applied force for simulated textures (this is scheduled).
- Experiments to test the various hypotheses for the Tardis Effect.
- Check the hypothesised fundamental parameter of roughness quantitatively and design experiments to test & refine it.
- Confirm that the results are algorithm-independent as well as hardware independent.

Other experiments particularly important for BT include ones on the requirements for networked haptics & multimodal interactions between haptics & the already transmitted audio-visual senses.

11. CONCLUSIONS

The sense of touch has become the third element of computer multimedia following sight & sound. This is especially useful for blind users who were being excluded by the growth of naïvely designed popular user computer desktop & World Wide Web displays with purely graphical controls but is also helpful to sighted users for whom it both aids normal use and enables totally new applications. However, it has implications for telephone companies providing computer networks. Even the fundamentals of what exactly needs to be stored, transmitted & output to cope with feelable computer systems were unknown. Questions to be answered included: “What exactly is it that gives the sensation of roughness?”; “Are there equivalents of visual illusions that will need compensation?”; & “Do there need to be settings that are adjustable between different users?”. This dissertation reported research that addressed these & other questions.

The fundamental basis of the feeling of roughness was traced to the magnitude of the stick-slip friction as the surface is traversed. Several illusions were discovered including a strong repeatable one named ‘The Tardis Effect’ whereby objects felt from inside feel bigger than when felt from outside. A parameter akin to visual-display gamma correction was proposed as one setting that may need to be adjustable for textures to feel the same to different users.

These discoveries will aid BT meeting the forthcoming need for support of haptic interaction across its networks.

12. REFERENCES

The abbreviation 'ASME-DSC' used in several references below refers to the Proceedings of the ASME's Dynamic Systems & Control Division at the ASME's International Mechanical Engineering Congress & Exposition (ASME is the American Society of Mechanical Engineering). The haptics subsection of this conference has become the world's main annual meeting of people working in computational haptics.

The few references (5 out of 87) which are to ephemeral & unrefereed WWW pages are individually justified.

J.L.Alt, **D.I.Rigas**, 1998, Communicating Graphical Information to Blind Users Using Music: The Role of Context in Design; Proc. CHI 98, 574-81.

S.Appelle, **F.J.Gravetter**, **P.W.Davidson**, 1980, Proportion Judgements in Haptic and Visual Form Perception; Canadian J. Psychology, 34, pp. 161-74.

P.J.Berkelman, **R.L.Hollis**, 1998, Haptic Interaction using Magnetic Levitation; ASME-DSC, 64, pp. 187-8.

P.J.Berkelman, **Z.J.Butler** **R.L.Hollis**, 1996, Design of a Hemispherical Magnetic Levitation Haptic Interface; ASME-DSC, 58, pp. 483-8.

K.R.Boff, **J.E.Lincoln**, 1988, Haptic Perception of Curvature: Effect of Curve Orientation & Type of Arm Movement; Engineering Data Compendium: Human Perception & Performance, AAMRL, pub. Wright-Patterson, section 6.609, pp. 1358-9.

B.Bongers, 1997, Tactile Display In Electronic Musical Instruments; Digest of IEE Comp. & Control Div. colloquium on Developments in Tactile Displays, 1997-01-21, Digest No. 96/012, pp. 7/1-3.

M.Brady, 2000; BSc (Hons) final year project dissertation, Dept. Psychology, Univ. Herts., UK.

F.H.Brigham, 1989, Statistical Methods for Testing the Conformance of Products to User Performance Standards; Behaviour Inf. Tech., 8(4), pp. 279-83.

S.Brewster, **M.Montgomery-Masters**, **A.Glendye**, **N.Kriz**, **S.Reid**, 2000, Haptic Feedback in the Training of Veterinary Students; downloaded 2000-06-13 from

<http://www.dcs.gla.ac.uk/~stephen/vets/>. (This described a use of haptics which is summarised in this dissertation. Unavailability of the original publication does not invalidate the potential uses.)

S.Brewster, H.Pengelly, 2000, Visual Impairment, Virtual Reality & Visualisation; downloaded 2000-06-13 from <http://www.dcs.gla.ac.uk/~stephen/visualisation/> . (As previous.)

T.Bruns, 1998, Feeling Nothing: Haptic Object Recognition with Blind and Sighted People in a Virtual Environment; P16, BSc (Hons) final year project dissertation, Dept. Psychology, Univ. Herts., UK.

P.Buttolo, R.Obao, B.Hannaford, W.McNeely, 1996, Force Feedback in Shared Virtual Simulations, Proc. MICAD, Paris.

Z-L.Cai, J.Dill, S.Paryandeh, 1999, Haptic Rendering: Practical Modelling and Collision Detection; ASME-DSC, 67, pp. 81-6.

C.G.L.Cao, C.L.MacKenzie, S.Payandeh, 1996, Task and Motion Analyses in Endoscopic Surgery; ASME-DSC, 58, pp. 583-90.

P-L.Chau, A.J.Hardwick, 1998, A New Order Parameter for Tetrahedral Configurations; Molecular Physics, 93(3), pp. 511-8.

C.E.Sherrick, R.W.Cholewiak, 1986, Cutaneous Sensitivity; Handbook of Perception and Human Performance (Ed. L.Bof, L.Kaufman, J.P.Thomas), pub. Wiley, 1(3) chapter 12.

F.J.Clark, K.W.Horch, 1986, Kinesthesia; Handbook of Perception and Human Performance (Ed. L.Bof, L.Kaufman, J.P.Thomas), pub. Wiley, 1(3) chapter 13.

C.Colwell, H.Petrie, D.Kornbrot, A.Hardwick, S.Furner, 1998a, Use of a Haptic Device by Blind and Sighted People: Perception of Virtual Textures and Objects; submitted to CHI'98.

C.Colwell, H.Petrie, D.Kornbrot, A.Hardwick, S.Furner, 1998b, Haptic Virtual Reality for Blind Computer Users; Proc. Third Int. ACM Conf. on Assistive Technologies (ASSETS '98), ACM Press.

F.L.Engel, P.Goossens, R.Haakma, 1994, Improved Efficiency through I- & E-feedback: a Trackball with Contextual Force Feedback; Int. J. Human-Computer Studies, 41, 949-74.

J.Ekberg, 1999, Cost 219 bis (Telecommunications: Access for Disabled People & The Elderly), 2000-3-14. Downloaded 2000-06-14 from <http://www.stakes.fi/cost219/> . (COST is a framework for scientific and technical co-operation in Europe from 25 member countries, see <http://www.belspo.be/cost/> , of which 219 bis is one of their projects).

Y.Fukui, M.Shimojo, 1992, Difference in Recognition of Optical Illusion Using Visual and Tactual Sense; *J. Robotics and Mechatronics*, 4(1), pp. 58-62.

S.Furner, 1998, May the Force be with You; *Ability* 26, pp. 8-9.

S.Furner, A.Hardwick, C.Colwell, T.Bruns, H.Petrie, D.Kornbrot, 1999 Computer Haptics - Putting Physical Sensation into the "Look and Feel" of Human Computer Interaction at the Desktop; *BT Technology Journal*, 17(1), accepted for publication but omitted from final print for editorial reasons.

R.B.Gillespie, M.S.O'Modhrain, P.Tang, D.Zaretsky, C.Pham, 1998, The Virtual Teacher; *ASME-DSC*, 64, pp. 171-8.

M.J.Gione, G.C.Burdea, M.Bouzit, 1999, The "Rutgers Angle" Orthopaedic Rehabilitation Interface; *ASME-DSC*, 67, pp. 305-12.

D.F.Green, J.K.Salisbury, 1997, Texture Sensing and Simulation using the PHANToM: Towards Remote Sensing of Soil properties; *Proc. 2nd PHANToM Users Group Workshop* (Ed. J.K.Salisbury & N.A.Srinivasan), A.I. Tech. Report 1617 & R.L.E. Tech. Report 618, MIT.

N.Hammond, J.Morton, A.MacLean, P.Barnard, 1983, Fragments and Signposts: User's Models of the System; *Proc. 10th Int. Symp. on Human Factors in Telecommunication*, pp 81-88.

B.Hannaford, S.Venema, 1995, Kinaesthetic Displays for Remote and Virtual Environments; *Virtual Environments and Advanced Interface Design*, Ed W.Barfield & T.Burness, pub. Oxford Univ. Press, pp. 415-470.

A.J.Hardwick, 2000, Haptic Simulation for Psychophysical Investigations; MSc dissertation, BT / Univ. London.

A.J.Hardwick, 1995, The Mechanism of Subharmonic Ultrasound Modulation by Forcibly Oscillated Bubbles; *Ultrasonics*, 33(4), pp. 341-3.

A.J.Hardwick, A.J.Walton, 1994, Forced Oscillations of Bubble in a Liquid; European J. Physics, 15, pp. 325-8.

A.J.Hardwick, A.J.Walton, 1995, The Acoustic Bubble Capacitor: a New method for Sizing Gas Bubbles in Liquids; Measurement Sci. & Tech., 6, pp. 202-5.

A.Hardwick, S.Furner, J.Rush, 1997b, Tactile Access for Blind People to Virtual Reality on the World Wide Web; IEE Digest No. 96/012 (Developments in Tactile Displays Colloquium), pp. 9/1-9/3.

A.Hardwick, S.Furner, J.Rush, 1998, Tactile Display of Virtual Reality from the World Wide Web - a Potential Access Method for Blind People; IEE Displays 18, pp. 153-61.

A.Hardwick, J.Rush, S.Furner, J.Seton, 1997a, Feeling It as well as Seeing It - Haptic Display within Gestural HCI for Multimedia Telematics Services; Progress in Gestural Interaction (Proc. Gesture Workshop '96), Ed. P.A.Harling & A.D.N.Edwards, pub. Springer-Verlag, ISBN 3-540-76094-6, pp. 105-116.

S.Hasegawa, M.Ishii, Y.Koike, M.Sato, 1999, Inter-process Communication for Force Display of Dynamic Virtual World [sic]; ASME-DSC, 67, pp. 211-8.

C.J.Hasser, A.S.Goldenberg, K.M.Martin, L.B.Rosenberg, 1998, User Performance in a GUI Pointing Task with a Low-cost Force-Feedback Computer Mouse; ASME-DSC, 64, pp. 151-6.

C.M.Hendriz, P-M.Cheng, W.K.Durfee, 1999, Relative Influence of Sensory Cues in a Multimodal Virtual Environment; ASME-DSC, 67, pp. 59-64.

C.Ho, C.Basdogan, M.Slater, N.Durlach, M.Srinivasan, 1998, The Influence of Haptic Communication on the Sense of Being Together; Proc. Int. Workshop on Presence in Shared Virtual Environments, BT Labs, June 1998.

C.Ho, C.Basdogan, M.Srinivasan, 1999, Efficient Point-Based Rendering Techniques for Haptic Display of Virtual Objects; Presence, 8(5), pp. 477-91.

Y.Ikei, M.Yamada, S.Fukuda, 1999, Tactile Texture Presentation by Vibratory Pin Arrays Based on Surface Height Maps; ASME-DSC, 67, pp. 51-8.

F.Infed S.V.Brown, C.D.Lee, D.A.Lawrence, A.M.Dougherty, L.Y.Pao, 1999, Combined Visual/Haptic Rendering Modes for Scientific Visualisation; ASME-DSC, 67, pp. 93-100.

G.Jansson, H.Petrie, C.Colwell, D.Kornbrot, J.Fänger, H.König, K.Billberger, A.Hardwick, S.Furner, 1999, Haptic Virtual Environments for Blind People: Exploratory Experiments with Two Devices; *Int. J. Virtual Reality*, 4(1), pp. 10-20.

R.L.Klatzky, S.J.Lederman, C.Hamilton, G.Ramsay, 1999, Perceiving Roughness via a Rigid Probe: Effect of Exploration Speed; *ASME-DSC*, 67, pp. 29-33.

S.Lakatos, L.E.Marks, 1998, Haptic Underestimation of Angular Extent; *Perception*, 27(6), pp. 737-54.

S.J.Lederman, 1974, Tactile Roughness of Grooved Surfaces: the Touching Process and Effects of Macro- and Microsurface Structure; *Perception & Psychophysics* 16(2), pp. 385-395.

S.J.Lederman, 1981, Perception of Surface Roughness by Active and Passive Touch; *Bulletin of the Psychonomic Soc.* 18(5), pp. 253-255.

S.J.Lederman, R.L.Klatzky, 1987, Hand Movements: A Window into Haptic Object Recognition; *Cognitive Psychology*, 19, pp. 342-368.

S.J.Lederman, R.L.Klatzky, 1998, Feeling through a Probe; *ASME-DSC*, 64, pp. 127-31.

S.J.Lederman, M.M.Taylor, 1972, Fingertip Force, Surface Geometry, and the Perception of Roughness by Active Touch. *Perception and Psychophysics*, 12(5), pp. 401-408.

J.M.Loomis, S.J.Lederman, 1986, Tactual Perception; *Handbook of Perception and Human Performance* (Ed. L.Bof, L.Kaufman, J.P.Thomas), pub. Wiley, 2(5) chapter 31.

W.L.Bryan, N.Harter, 1897, Studies in the Physiology and Psychology of the Telegraph Language; *Psychological Rev.*, 4, pp. 27-53.

J.McCrone, 1993, *The Myth of Rationality*; pub. Macmillan, ISBN 0-333-57284x, p. 57.

L.Parkin, 1996, *Doctor Who: A History of the Universe*; pub. Doctor Who Books (a Virgin Pub. Ltd imprint), ISBN 0-426-20471-9.

E.Pere, D.Gomez, G.Burdea, N.Langrana, 1996, PC-based Virtual Reality System with Dextrous Force Feedback; *ASME-DSC*, 58, pp. 495-502.

K.E.MacLean, 1996, The "Haptic Camera": A Technique for Characterising and Playing Back haptic Properties of Real Enviroments; ASME-DSC, 58, pp. 459-67.

T.H.Massie, 1996, Initial Haptic Explorations with the PHANToM: Virtual Touch through Point Interaction; Master's thesis, MIT.

J.Mathews, R.L.Walker, 1970, Mathematical Methods of Physics; Second Edition, pub. Addison-Wesley, ISBN 0-8053-7002-1.

B.E.Miller, J.E.Colgate, 1998, Using a Wavelet Network to Characterise Real Environments for Haptic Display; ASME-DSC, 64, pp. 257-64.

B.E.Miller, J.E.Colgate, R.A.Freeman, 1999, Computational Delay and Free Mode Environment Design for Haptic Display; ASME-DSC, 67, pp. 229-36.

M.D.R.Minsky, 1995, Computational Haptics: The Sandpaper System for Synthesizing Texture for a Force-Feedback Display; PhD Thesis, Massachusetts Inst. Tech..

D.Minsky, S.J.Lederman, 1996, Simulated Haptic Textures: Roughness; ASME-DSC, 58, pp. 421-6.

H.B.Morgenbesser, M.A.Srinivasan, 1996, Force Shading for Haptic Shape Perception; ASME-DSC, 58, pp. 407-12.

A.Murray, R.L.Klatzky, P.K.Khosda, 1999, Summation Multifinger Vibrotactile Stimuli; ASME-DSC, 67, pp. 1-8.

T.Nara, T.Maeda, Y.Yanagida, S.Tachi, 1999, A Tactile Display Using Ultrasonic Elastic Waves in a Metal Tapered Membrane; ASME-DSC, 67, pp. 283-8.

M.Ottensmeyer, 1997, Developing a "thermal display"; downloaded 2000-06-13 from <http://www.mit.edu/people/markott/haptics.html> . (This described a physical output method which was summarised in this dissertation. Unavailability of the original publication does not invalidate its potential as an output method.)

J.Payette, V.Hayward, C.Ramstein, D.Bergeron, 1996, Evaluations of a Force Feedback (Haptic) Computer Pointing Device in Zero Gravity; ASME-DSC, 58, pp 547-53.

D.K.Pai, L-M.Reissell, 1996, Touching Multiresolution Curves; ASME-DSC, 58, pp. 427-32.

D.T.V.Pawluk, C.P.vanBuskirk, J.H.Killebrew, S.S.Hsiao, K.O.Johnson, 1998, Control and Pattern Specification for a High Density Tactile Array; ASME-DSC, 64, pp. 97-101.

H.Petrie, 1997, User-Centred Design and Evaluation of Adaptive and Assistive Technology for Disabled and Elderly Users; Informationstechnik und Technische Informatik: IT + TI, 39, pp. 7-12.

P.Penn, 2000, Haptic Perception in Virtual Reality: the Impact of Visual Status and Device Parameters on the Perception of Texture and Objects; MPhil to PhD Progression Report, Sensory Disabilities Res. Unit, Dept. Psycho., Univ. Hertfordshire, UK.

PC Pro, 1997, Teletouch the Web; PC Pro, November,.

C.Ramstein, 1996, Combining Haptic and Braille Technologies: Design Issues and Pilot Study; ASSETS '96 (Vancouver, Canada), pp. 37-44.

G.Revesz, 1950, Psychology and Art of the Blind; pub. Longmans.

M.Shibita, R.D.Howe, 1999, The Effects of Gloves on the Performance of a Tactile Perception Task and Precision Grasping; ASME-DSC, 67, pp. 9-17.

M.A.Srinivasan, C.Basdogan, 1997, Haptics in Virtual Environments: Taxonomy, Research Status & Challenges; Computers & Graphics, 21(4), (Editorial overview to special issue).

S.S.Stevens, J.R.Harris, 1962, The Scaling of Subjective Roughness and Smoothness. J. Exptl. Psychology, 64, pp. 489 - 494.

D.J.Sturman, D.Zelter, 1994, A Survey of Glove-based Input; IEEE Computer Graphics & Applications, Jan, pp. 30-8.

P.M.Taylor, A.Moser, A.Creed, 1997a, The Design and Control of a Tactile Display based on Shape Memory Alloys; Digest of IEE Comp. & Control Div. colloquium on Developments in Tactile, 1997-01-21, Digest No. 96/012, pp. 1/1-4.

P.M.Taylor, A.Hosseini-Sianaki, C.J.Varley, D.M.Pollet, 1997b, Advances in an Electrorheological Fluid Based Tactile Array; Digest of IEE Comp. & Control Div. colloquium on Developments in Tactile, 1997-01-21, Digest No. 96/012, pp. 5/1-5.

T.V.Thompson.II, E.Cohen, 1999, Direct haptic Rendering of Complex

Trimmed Nurbs Models; ASME-DSC, 67, pp. 109-116.

L.Vaas, 2000, Web's Blind Spot: Disabled Users; ZDNet News, Mon 2000-4-17, downloaded 2000-06-13 from <http://www.zdnet.co.uk/news/2000/15/ns-14860.html> . (This is newspaper article. ZDNet news is one of the most popular serious computer industry newspapers but is only available on-line. I have found its archive to have been reliable so far.)

R.A.Virzi, 1992, Refining the Test Phase of Usability Evaluation: How Many Subjects is Enough?; Human Factors, 34(4), pp. 475-68.

S.A.Wall, W.S.Harwin, 1999, Modelling of Surface Identifying Characteristics using Fourier Series; ASME-DSC, 67, pp. 65-71.

R.Waller, 1995; personal communication from Bob Waller, BT.

J.M.Weisenberger, M.J.Krier, M.A.Rinker, S.M.Kreidler, 1999, The Role of the End-effector in th the Perception of Virtual Surfaces Presented via a Force-feedback Haptic Probe; ASME-DSC, 67, pp. 35-41.

R.L.WilliamnsII, 1998, Cable-suspended Haptic Interface; ASME-DSC, 64, pp. 207-12.

13. APPENDICES

13.1. APPENDIX A: KINEMATICS EQUATIONS FOR IE 3000

The IE 3000 has a mechanical linkage that is equivalent to that in Figure 33. The drivers that were provided by Immersion were seriously faulty in that they assumed the Cartesian co-ordinates of the probe tip varied linearly & independently with the angular readings from encoders on the three motors. This was not really the case & the misassumption caused severe distortion. For example a cube so represented has a convex front face, concave back face & sides slope in backwards. The correct co-ordinate transform needed to be calculated.




Figure 33: Schematic IE 3000 linkage.

The mechanical co-ordinates were defined follows: α = the angle of rotation about the vertical axis; β = the angle of rotation about the horizontal axis; L = the distance of the end of the rod from the intersection point of the two axes; and the rotational origins are defined so that the rod is centralised pointing horizontally towards the user (i.e. $x = y = 0$) when $\alpha = \beta = 0$. The Cartesian co-ordinates of the probe tip can be related to these by considering intersections of the two sets of

ellipses in the projections of the loci of the probe tip on the x-y plane for the cases of constant α with varying β and constant β with varying α . The resulting relations are*

$$x = LA^{-1/2} \sin \alpha \cos \beta - (\text{x axis Origin}) \quad \{4\}$$

$$y = LA^{-1/2} \cos \alpha \sin \beta - (\text{y axis Origin}) \quad \{5\}$$

$$z = LA^{-1/2} \cos \alpha \cos \beta - (\text{z axis Origin}) \quad \{6\}$$

$$\text{where } A = 1 - \sin^2 \alpha \sin^2 \beta.$$

A kinematic solution for force feedback not only needs a position conversion for input but a force conversion for output. Whereas the co-ordinates needed to be converted from the mechanical (α, β, L) system to the Cartesian (x, y, z) system, forces needed to be converted from Cartesian components $(F_x, F_y$ and $F_z)$ to the mechanical torques $(T_\alpha$ and T_β about the α and β axes respectively) and extension force $(F_L$ along the L axis) required from the motors. To do this, the three basis vectors in the mechanical system, $\hat{\alpha}$, $\hat{\beta}$ and \hat{L} , were expressed in terms of the Cartesian basis vectors, \hat{x} , \hat{y} and \hat{z} , by partially differentiating the general Cartesian expression for a position, $x\hat{x} + y\hat{y} + z\hat{z}$, with respect to α , β and L individually after substituting in x , y and z from equations 4 to 6. The general Cartesian expression for a force, $F_x\hat{x} + F_y\hat{y} + F_z\hat{z}$, could then set equal to the mechanical one, $(T_\alpha/L)\hat{\alpha} + (T_\beta/L)\hat{\beta} + F_L\hat{L}$. Comparing the coefficients of the basis vectors then gave the conversion from Cartesian force components to the torques and force required from the motors. The resulting conversion was simply

$$T_\alpha = LA^{3/2} \left(\frac{F_x \cos \alpha - F_z \sin \beta}{(\cos^2 \alpha + \sin^2 \alpha \cos^2 \beta) \cos \beta} \right) \quad \{7\}$$

$$T_\beta = LA^{3/2} \left(\frac{F_y \cos \beta - F_z \sin \alpha}{(\cos^2 \beta + \sin^2 \beta \cos^2 \alpha) \cos \alpha} \right) \quad \{8\}$$

$$F_L = LA^{1/2} \left(\frac{\frac{F_x \cos \beta \sin \alpha}{\cos^2 \alpha + \sin^2 \alpha \cos^2 \beta} + \frac{F_y \cos \alpha \sin \beta}{\cos^2 \beta + \sin^2 \beta \cos^2 \alpha} + \frac{F_z A \cos \alpha \cos \beta}{(\cos^2 \alpha + \sin^2 \alpha \cos^2 \beta)(\cos^2 \beta + \sin^2 \beta \cos^2 \alpha)}} \right) \quad \{9\}$$

* Due to a typographical error, equations 4 to 6 were included in an earlier paper [A.Hardwick *et al* 1997a] with the $A^{-1/2}$ normalisation factors omitted.

The haptic simulation software continually cycles and on each cycle it inputs α , β and L from the probe's encoders, applies equations 4 to 6 to convert them to x , y and z , calculates the force components, F_x , F_y and F_z , required at that position, applies equations 7 to 9 to convert them to T_α , T_β and F_L , and outputs them to the probe's motors.

13.2. APPENDIX B: CLASS STRUCTURE

The different primitive haptic shapes naturally fit within an object-orientated scheme for the computer programming. A base-class for haptic objects was defined and equipped with member functions to carry out the manipulations that are common to all simulated solids such as rotations and translations. From this base-class separate classes were derived for each of the primitive objects. Three virtual functions had to be defined for each primitive object: one that allowed instances of the class to be copied and duplicated; one that determined if a given point was within the shape; and one which calculated the force vector at that point. Functions to allow the variable parameters of the shapes (radii, textures, skin thicknesses etc.) were also added. Because the functions that are used for determining the forces are virtual, only the parts of the program which set up the haptic scene need to know about the different shapes. The rest of the program just used calls to functions in the haptic object base class of the instances and could ignore the fact that totally different algorithms may be actually being called depending upon the particular derived class of which it is an instance. The haptic simulation loop only knows that it has been given an object to simulate and that the object can work out if a given point is in it and, if it is, what force ought to be output; it does not need to know anything else about the object.

Texture classes are similarly based on a single haptic texture base class. The virtual functions which are defined for each derived texture class are the copying function and one that returns a height for a given point on a plane. Each haptic object can have a haptic texture (or more than one if the object has multiple faces) but only needs to know that it is something which returns a height as a function of position, nothing more. Hence extra textures can be created in the future without rewriting the haptic object simulation functions. The class hierarchy is shown schematically in Figure 34.

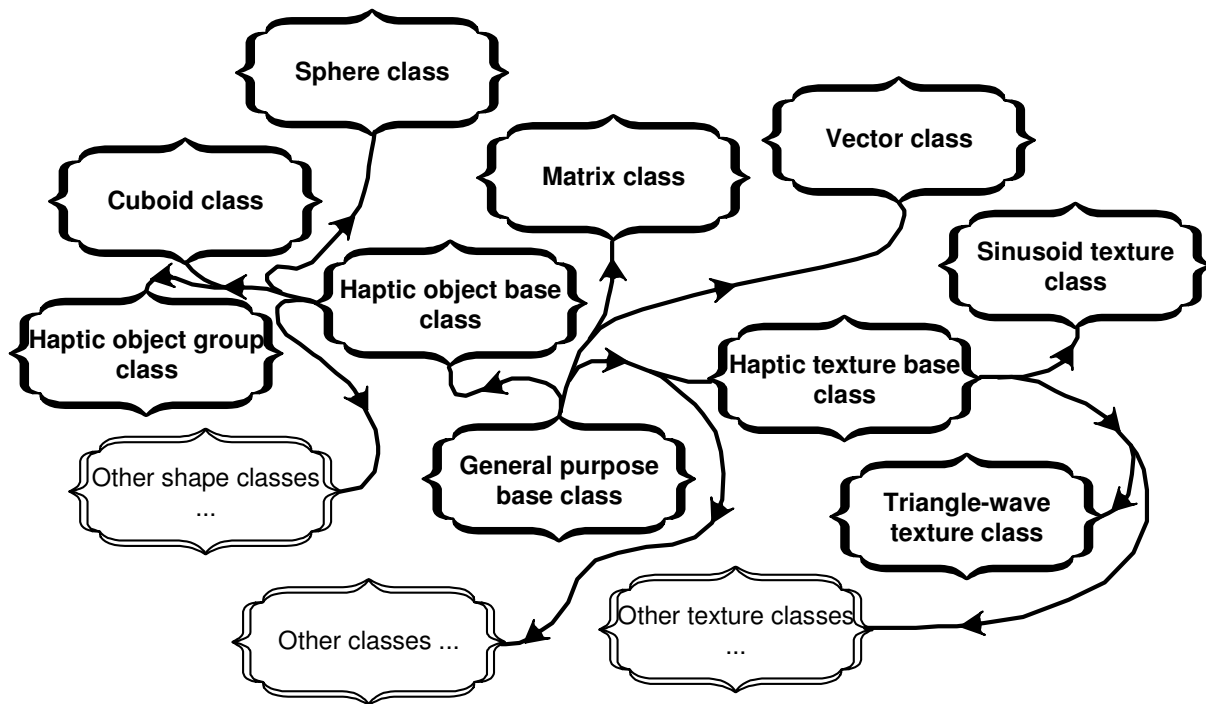


Figure 34: The class structure of the haptic objects.

13.3. APPENDIX C: IMPULSE ENGINE 3000 DESIGN FAULTS

Many faults were found in the design of the IE 3000 & it frequently suffered mechanical & electrical failures. Several are listed below lest as advice for future haptic hardware designers. These have been fed back to the manufacturers &, where possible, the device was modified by ourselves to correct for the them.

- Solid core ribbon cable was used in a flexing situation. It broke from metal fatigue.
- No external plugs, just a fixed ribbon cable, between the ISA board in the PC and the amplifier box. This made transportation difficult.
- There was no bracket on the ISA board. This allowed the board to slip partly out the socket & blow.
- ISA board's addressing system ignored several address lines and therefore aliased itself throughout the addressing space & so it was unnecessarily likely to conflict with other devices.
- Excessive slack in the stylus universal joint reduced the finesse of the device greatly.
- Thin weak bolts were used as pivots. They sheared off quickly.
- Lack of ball bearing races on the pivots made the pivoting rough and quickened wear.
- Drive cord around the motor spindle was arranged so that it chafed against itself & wore though.
- Slider could twist on its axis and so jam. Inexperienced users forcing through this damaged the device.

13.4. APPENDIX D: HAPTIC NETWORKING ISSUES

13.4.1. INTRODUCTION & CLASSIFICATION

This appendix provides a brief overview of the major issues in haptic networking.

End-to-end transmission of real haptic information, rather than merely using haptics to enhance user interfaces that are not fundamentally haptic, presents some interesting additional challenges for network technology and user interface designing that are obviously relevant to telecommunications companies. It will be useful here to classify haptic transmission systems into 3 categories according to the type of physical environment at the remote end that needs to be haptically transmitted. For simplicity the output is assumed to be by force-feedback.

- **Static or fully predictable:** In this case a model of the environment can be built up and batch transmitted. All further processing is at the receiving end of the link where a pure simulation algorithm can be used. The only extra difficulty beyond that of local simulation is building up the haptic model at the transmitting end from automated investigation of the environment.
- **Mechanical & not fully predictable:** A model of the system cannot be constructed in advance because the behaviour of the environment is either non-deterministic or dependent on the user's actions in an irreversible manner so the responses cannot be tested. This needs haptic transmission in both directions so that the forces &/or motions that the user makes must be mimicked by a slave device at the other end of the link. The slave device then, in moving, interacts with its environments & the forces it experiences from contact with real objects are sent user end of the link for output to the user.
- **Human:** With humans at both ends of the link the situation becomes an extreme case of the mechanical unpredictable one. Humans are very unpredictable machines. The transmission may also need to be fully symmetric so that the persons can feel each other.

The first of the three is just requires a network over which a data file can be sent but the other two categories require real-time duplex transmission so bandwidth, latency & reliability become crucial factors.

13.4.2. NETWORK 'BANDWIDTH' PROBLEMS

Data rate* limitation is not a serious problem for a point-contact haptic transmission. A 16-bit resolution 3d force vector sent 1000 times per second (as is output to a Phantom) only requires about 47 Kbit/s. This could easily be reduced by differential encoding & relative scaling. The reason for the low data rate compared to visual transmission is that although the frame rate is much higher, each frame is only equivalent to a single pixel picture. If hardware improves so that a full 2d surface contact can be accurately simulated, not just a single point, the required data rate may be much higher.

The data could also be decimated and the missing values replaced by linear extrapolation at the receiver since the scene being felt varies very slowly compared to the update frequency needed in the force-feedback loop between the user's hand & the simulation hardware to ensure stability. This extrapolation is effectively a synchronised local-model scheme. The local force-feedback loop handles immediate changes in forces due to small motions of the local user's hand based using its local model & that model is periodically updated to match changes in the remote environment. Linear extrapolation is a very simple model which still requires moderately frequent updates to ensure that abrupt spatial changes in force - such as from hitting a hard surface - do not cause instability in the simulation. A better extrapolation scheme using predictive or model-based encoding to remove regular features from the transmitted data stream could substantially reduce the required data rate. The minimum data rate needed will, of course, depend on the nature of the haptic interaction required.

13.4.3. NETWORK LATENCY PROBLEMS

Latency is a considerably more serious problem for networked haptic interaction than is bandwidth. A good haptic force-feedback simulation requires updates every 10 milliseconds or so but the delays involved in modern data

* In digital transmission usage, 'data rate' is often referred to as 'bandwidth' & 'bit rate' as 'baud rate' copying the terms from analogue transmission usage. However, the two sets are not necessarily the same. 'Baud rate' means 'symbols per second' not 'bits per second' which is different unless a simple 2-state modulation scheme like Binary Modulation or Binary Phase Shift Keying is used. Other schemes make more efficient use of bandwidth by multilevel amplitude modulation & separate modulation of quadrature carrier components so that each symbol transmitted represents more than one bit so the bit rate is higher than the baud rate & the data rate is higher than the bandwidth. The assumption that the two sets of terms are equal probably originates from computer scientists being familiar with networks based on IE802 Ethernet, IE802 token ring & fibre optics which use 2-state modulation for easy noise-rejection. In telecommunications, bandwidth is limited by the properties of existing copper pair cables in local access wiring or by the allocation of radio spectral bands. Therefore multilevel & spread-spectrum modulations are used. This applies not only to new technologies like ADSL but even to old ISDN which uses 2B1Q (2 bits per 1 quad, i.e. 4-level) modulation.

networks are usually much higher than this. Obviously, one answer is to use a connection-orientated network like PSTN or ISDN instead of a packet based network like TCP/IP or Ethernet. This will avoid the packetisation delays and roundabout high hop-count routing. Even so, the digital switching and multiplexing can cause some delay and the fundamental physical limit of the speed of light at 300km/ms is unavoidable.

For long distances and high latency transmission paths, it is necessary to find some way to avoid having to perform the actual force-feedback directly across the network. This essentially implies that a local model of some sort must be used. Even with a human at both ends, some limited amount of modelling can be made. Human limbs can be modelled as mechanical linkages with inertia, elasticity & damping and even with predictable reflex responses on timescales less than those which it takes for the higher centres of the brain to instruct muscles to perform voluntary movements.

Problems still remain if the environment can change unpredictably in less time than the round-trip latency. Systems used in the past for remote haptic interaction studies have employed a variety of methods to avoid actually networking unpredictable human contact across high latency links. The simplest method is to keep network distance very short; the studies of [Ho *et al* 1998] with 2 networked Phantoms actually had the Phantoms in neighbouring rooms wired to the same PC. The network length could be said to have been only a few centimetres on a PCI bus. Another is turn-taking where only one user at a time has a controlling influence on certain objects and only that user has a haptic response from that object. An example of this was a networked squash game in which players took turns at hitting the ball [Buttolo *et al* 1996]. A third is to add a large amount of damping, inertia or slack in the simulation to hide delay problems. Mediating the touch between the users through a virtual object can do this. For example, if the users simultaneously manipulate the same object by different point contacts then the inertia & twisting of object can suffice. This was also used in a trial of pairs of users in different locations attempting to co-operatively move a simulated ring along a simulated wire [Ho *et al* 1998].

13.4.4. OUR HAPTIC NETWORKING SYSTEM & EXPERIMENTS

Our haptic displays work has not only created psychological experiment & demonstration systems but also a haptic network system. This linked 2 Phantoms over a TCP/IP network. It allowed either 2 remote people to interact, one using each Phantom, or a user to feel a remote physical environment, by using one Phantom as

a remote slave.

It was written in a mixture of C++ & LabView & functioned by transmitting the position of each Phantom handpiece to the other via UDP packets. Its operation was fully symmetric even when running as a single person remote feeling device. The algorithms used for local force-feedback included several obvious robustness measures. Firstly, all packets were sequence numbered and any found to arrive out of order were discarded. Secondly, it never waited for packets but in the absence of a new packet with updated information arriving it simply approximated the current remote position to be equal to last received position (an improved version that estimated the position change using velocity & acceleration calculated from past position history was intended but not found to be vital). Thirdly, it did not try to reproduce the real live force changes from the other end but simply to spring to its current relative position. This third measure is simple but its implications are rather profound. It is effectively an extremely simple local model which means that the force-feedback loop is virtually a 1 kHz local one regardless of the network characteristics. It also converted displacement-feedback to force-feedback. The networking was really displacement-feedback not force-feedback; force-feedback only functioned in the local models at both ends of the network.

No formal experiments have yet been performed but this simple system was subjectively found to be satisfactory running on the BT laboratory intranet. Experimenters & visitors who tried it were able to move each other's movements*. They also felt the shapes and deformabilities of real objects. Somewhat unexpectedly, it was found that decimating the data down from 1000 updates per second to just 50 still gave a satisfactory response. Indeed at more than 300 packets per second the Windows/LabView UDP reception overloaded resulting in an ever expanding incoming buffer; therefore decimation on sending was not just acceptable but vital. Another problem with the underlying UDP stack was that a latency of at least 20 ms was found even between PCs on the same Ethernet hub. However, that also had a beneficial consequence since it proved the system to be resilient against that degree of latency.

Usability studies are planned as is testing the system on the public Internet. Internet haptic transmission between BT Labs and the University of Hertfordshire

* For the record, the first people to interact across the system were Andrew Hardwick (who created it) & Paul Penn (who was writing a GHOST manual for psychologists in the same room). The first real object remotely felt was an 'aerosol' compressed air duster canister.

has been attempted but the BT firewall (which compares the addresses of incoming UDP packets to previous outgoing ones) added far too much delay. Future tests will use dial-up networking.

13.5. APPENDIX E: ACRONYMS

1d, 2d, 3d	1, 2, 3 dimensional
ADSL	Asymmetric Digital Subscriber Line
API	Application Programmers Interface
BT	British Telecommunications plc.
CHI	Computer Human Interaction
GUI	Graphical User Interface
HCI	Human Computer Interaction
HTML	Hypertext Mark-up Language
ISDN	Integrated Services Digital Network
MP3	MPEG 2 level 3 audio compression
MPEG	Motion Picture Experts Group
PCI	Peripheral Component Interconnect
RLSB	Royal London School for the Blind
RNIB	Royal National Institute for the Blind
SDRU	Sensory Disabilities Research Unit
TCP/IP	Transmission Control Protocol / Internet Protocol
UDP	User Datagram Protocol
VRML	Virtual Reality Modelling Language
WIMP	Windows Icons Menus Pointing-device or Windows Icons Mouse Pull-down-menus
WWW	World Wide Web

Table 10: Acronyms.

13.6. APPENDIX F: PUBLICATIONS FROM THIS WORK

Here is an overview of some publications from the Haptic Displays work*.

13.6.1. JOURNAL & CONFERENCE PAPERS

This work has generated 6 published papers & 2 near misses so far. The papers have covered both technology (those with Hardwick as first author) & psychology:

- **Hardwick et al 1997a:** Conference paper. Initial work with IE 3000. Presented, demonstrated & published in Progress in Gestural Interaction.

* It is very immodest to boast about ones own publications & other achievements but [redacted policy material] commercial environment & academic funding structures encourage, indeed enforce, publicity for advertising & for grant obtaining.

- **Hardwick et al 1997b:** Colloquium digest paper. Overview of simulation algorithms & initial IE 3000 observations. Presented & published in IE digest.
- **Hardwick et al 1998:** Refereed journal paper. Simulation algorithms & initial IE 3000 observations. Published in IEE Displays.
- **Furner 1998:** Requested paper. Overview for the general public. Published in BT Ability (disability magazine).
- **Colwell 1998b:** Refereed conference paper. Similar to previous. Accepted, presented & published at the Assets conference.
- **Furner et al 1999:** Refereed journal paper. Covered both algorithms & psychology for the BTTJ "Matching Technology to People" issue. Accepted by referees & sent to typesetters but omitted from final publication when the journal scope was reduced merely audio-visual.
- **Jansson et al 1999:** Refereed journal paper. A collaborative publication by the BT/SDRU team & a haptics group in the University of Uppsala in Sweden. Published in Int. J. Virtual Reality.

There are more papers scheduled based on the latest results [Penn 2000]. It will be possible to generate journal papers based on the haptic gamma correction & the roughness explanation from this report as well.

Articles have also been published by journalists about this work in many publications ranging from PC Pro [PC Pro 1997] to a local newspaper in Australia. The current author has also published papers on other subjects including both acoustic bubble measurement [Hardwick 1995, Hardwick & Walton 1994 & 1995] for the oil industry and mathematics [Chau & Hardwick 1998] for theoretical biochemistry.

13.6.2. WWW SITES

There are [*now "were"*] three WWW sites anent this project:

- **The BT internal site*** contains a comprehensive collection of information about the project (134 HTML pages & 227 images in 77 subdirectories) plus a collection of reprints to download (the papers plus 16 presentations).
- **The BT official public site[†]** overviews the project for public consumption.
- **The VVECC Site[‡]** was produced as a web version of the presentation given to VVECC.

* <http://www.amu.bt.co.uk/techinteg/projects/teletouch/index.html>

† <http://www.bt.com/innovation/exhibition/haptic/index.htm>

‡ <http://www-aisg.cis.rl.ac.uk/live/VVECC/proceed/perceptual/materials/hardwick/index.html>

These sites were produced by the current author and are blind-accessible except for the official public site which was produced by [redacted for diplomacy] (who design WWW sites as if they were paper publications & concentrate on initial appearance ignoring usability) based on the material in an earlier BT public site* by the current author which was blind-accessible.

13.6.3. PRESENTATIONS, VISITORS, SHOWS, DEMO'S ETC.

There have been far too many presentations & demonstrations of this work to report here. The following are few to show the variety of interest:

- **The London Royal Society for the Blind's Dorting College** was a demonstration location. It was very popular.
- **The Virtual Reality & Virtual Environments Community Club (VVECC)** requested a conference presentation at the Rutherford Appleton which was well received.
- **Open University** filmed the demo's for a course on evaluating user interfaces.
- **The head of Deutsche Telekom's Research Division** was amongst the first visitors to receive a demonstration of haptic networking.
- **Ipswich Engineering Society** were a local group who received a full evening seminar on the subject when the intended BT speaker on intelligent agents was unavailable & the current author stood in at short notice.
- **David Blunket**, the blind cabinet minister for education, tried the IE 3000 system at Northcote School in Wolverhampton (Figure 35).

* <http://www.labs.bt.com/projects/teletouch/index.htm> (no longer live)



Figure 35: David Blunket (right) feeling the solid object simulation [BT press photograph. Retouched by the current author.]

14. ACKNOWLEDGEMENTS

Thanks to: Stephen Furner for the BT part of the psychology, for handling BT's internal administration, for his work on the project publicity & his optimistic enthusiasm that goes with my pessimistic realism; Chetz Colwell, Paul Penn, Timo Bruns & Mark Brady for carrying out the laborious experiments; Helen Petrie for heading the SDRU team; Paul Penn for giving a psychologist's opinion of a draft copy; Brian Macdonald & Jim Alty for being the supervisors; & BT for funding it.