OCULUS INTOUCH

DESIGNING A CONTROLLER FOR COMPELLING VR EXPERIENCES THROUGH TOUCH

DANIEL SHOR

CONTENTS

CONTENTS (CONT'D)

FOREWORD

It's not the destination, but the journey. This old cliche has been my mantra through the adventure that has been this thesis. It has certainly been a journey. Over the past one hundred days, I have had the opportunity to meet and learn from some truly fantastic individuals. I have traveled to Berlin, Paris, Tokyo, and Seattle to sit down and learn from genuinely amazing and talented people willing to share their knowledge and enthusiasm with me. It is only with their help that I've been able to succeed.

I want to begin by thanking my chair, the wonderful Jess Hartcher-O'Brien. The completion of this thesis had its share of challenges, but the support, wisdom, and guidance I have received from her made it's success a foregone conclusion. Jess and I have been working together for almost two years, and every day, I'm astounded by the depths of knowledge and her willingness to devote time, energy, and effort on my behalf. Jess has given me space and encouragement to strike out in new directions, and the resources and network to push me forward. While I came to TU Delft interested in exploring haptics, without Jess, this would not be my path moving forward. My success is her success, and I am forever in her debt.

I am truly thankful for the level of trust, confidence, and support I received from my advisor, Henk Crone. His advice and suggestions come from immense knowledge and a truly unique ability to cut directly to the core of a subject. Without him, I would still be struggling to put a coherent direction and structure around the work I've done.

I'm also blessed by the resources provided to my thesis by the wonderful team at Oculus VR/Facebook Technologies. Robin Miller, Henric Jentz, and Mark Ando have devoted far more time that I could have reasonably expected to help my project along. Each has cleared out their calendar for weeks at a time and opened up their facilities and teams to help me succeed. My extended visits at the Oculus facility were enriching and instructive, and I can only hope this thesis begins to pay back all the support and knowledge they have sent my way. I'd also like to thank Jenny Spurlock and Khaled Borous and the rest of the truly wonderful people I met at the Oculus Stadium complex in Seattle. Their genuine curiosity and support were as energizing as the excellent coffee Seattle offers!

I'd also like to thank the incredible team at Actronika SAS -- Vincent Hayward, Rafal Pijewski, and Thomas Farnoux. Beyond tolerating my admittedly terrible French, the Actronika folks have gone above and beyond in helping me understand the complexities of actuator design and dropped everything to design and built some "incroyable" voice coil actuators.

My love and gratitude go out to the fantastic LoFelt GmbH team -- Daniel Büttner, Remi Laoubi, and Ben Schneiders. Beyond sending me as many LoFelt EVKs

as I could ever need, the LoFelt squad has stepped up time and time again to help me succeed.

Finally, I'd like to thank my friends and family for their love and warmth through the bright days and dark days of this thesis. The constant love and support coming in from nearly every time-zone in the world kept morale high.

To the TU Delft, the corridors of the faculty of Industrial Design Engineering, and my lovely desk at the Perceptual-Intelligence Touch lab I say: "Bedankt voor alles, en tot ziens!"

ABSTRACT / EXECUTIVE SUMMARY

Arthur C. Clarke said: "Any sufficiently advanced technology is indistinguishable from magic," and cutting edge Virtual Reality (VR) certainly is quite the illusion. VR opens to us a near infinite set of worlds and possibilities. Comprehensive, Immersive, Virtual Environments (IVEs) will change the way we shop, explore, love and live. They'll close the gap between our digital imaginations and our physical limitations.

The barrier to this vision is that right now, IVEs don't feel real. We may have breathtaking, high resolution displays, and compelling comprehensive spatial audio, but when we reach out and touch the digital world of an IVE, we're left holding nothing.

In this thesis, I propose a solution to this challenge, redesigning one of the most common VR ecosystems -- Oculus -- to contain cutting edge vibrotactile effects and actuators. The following pages lay out a narrative for a product -- the **Oculus inTouch** controller -- that balances performance, comfort, immersion and realism; yet remains accessible, flexible, and compelling.

More than a product, I introduce an approach -- **Interaction Centric Design** -- to take haptics from the intangible world of psychophysics, and the nuanced and intricate world of the engineer, placing it within the grasp of the designer. I define a new theory -- **Haptic Cuing** -- predicated on guiding user behavior through the careful selection of haptic signal to avoid producing haptic noise.

Over these next pages, I outline my journey into and through the world of haptics. I explain my process, my theory, and my path to designing what I hope is a foundational piece of haptic interfaces to come.

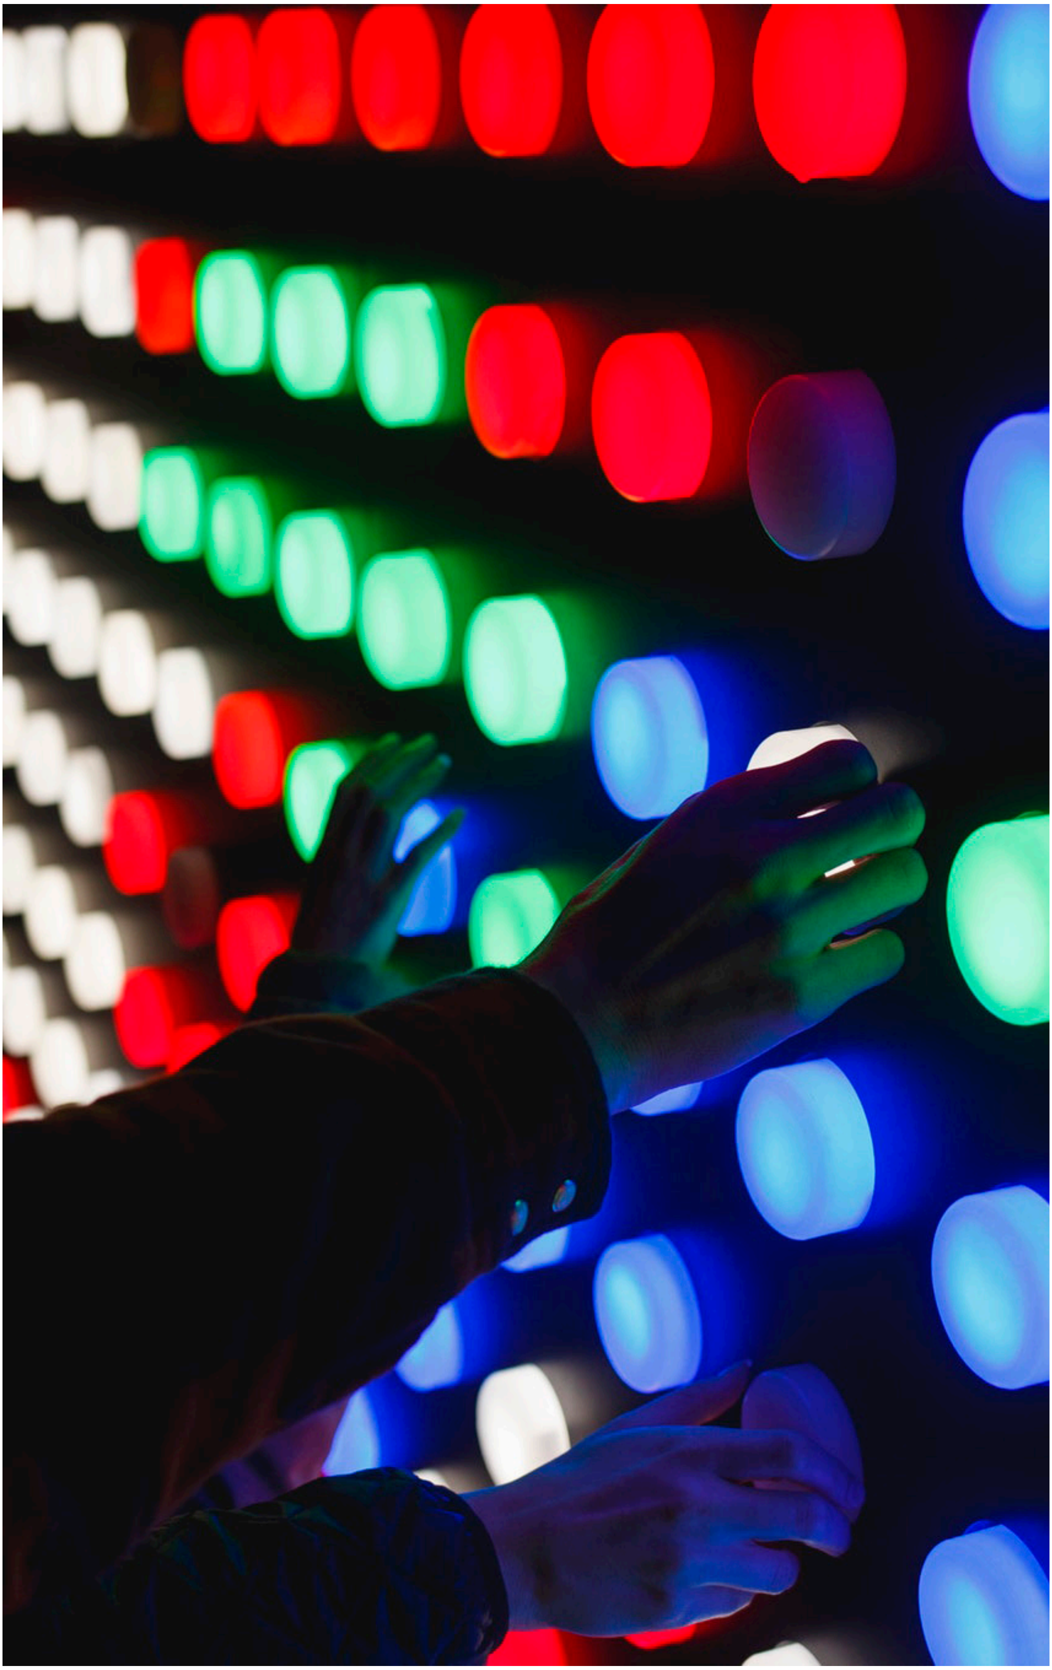

INTRODUCTION: A BACKGROUND TO HAPTICS AND VR

CHAPTER 01

INTRODUCTION: A BACKGROUND TO HAPTICS & VR

Virtual Reality (VR) is rapidly moving from novelty to everyday reality. Over the past decades, we have become familiar with a digital world. A world all around us that we access through two-dimensional windows -- the screens on our computers, phones, and smartwatches. This parallel reality is gaining depth, blurring the line between digital and physical. Increasingly, we shop, socialize, laugh, love, and live through the digital world.

VR aims to make this two-dimensional world three-dimensional. As we further erase the boundaries between the digital and the physical world, we will continue to change the way we work, play, and communicate. A significant player in the VR field, Facebook Technologies LLC "Oculus" produces VR hardware and software. The mission of Oculus is to have meaningful interactions regardless of physical distance. To achieve this mission, Oculus develops a family of Head-Mounted Displays (HMDs) and controllers for VR. The flagship line, the Oculus Rift/Oculus Rift-S requires a computer to render high-quality visuals and audio. Their Oculus Quest and Oculus Go lines are computer-free, relying on high-end mobile processors to create immersive, untethered VR experiences.

The core of virtual reality is an illusion -- the Immersive Virtual Environment (IVE). When a user dons an Oculus HMD, they are teleported into a construct -- a rendering of the physical world in the digital. This construct presents a computer-generated world as if it were real -- producing a sense of presence, or "being there," in the user's mind. It is possible to generate virtual environments with breath-taking graphics and vibrant, realistic acoustics. At their core, these renderings are always an intangible, digital, replication of the physical world -- it is impossible to reach out and touch the software.

VR adds another wrinkle into this challenge. While touchscreens require vibrotactile emulation of button effects, they can still rely on the skin displacement that stems from the contact between finger and screen. This is not the case in VR, where the user is interacting with purely digital surfaces without physical presence. These interactions violate some of the basic laws of physics -- when we press against a virtual surface, we don't experience a normal force.

Fig: Oculus Quest (Top) and Oculus Rift-S HMDs

THE OCULUS TOUCH & QUEST TOUCH CONTROLLER

To make these virtual worlds intractable, Oculus has a series of controllers branded under the Oculus Touch line. These controllers come in a handed set -- one left, one right -- each with a joystick, a pair of game buttons, a menu button, and two triggers. One trigger sits in the grip, and one on the front of the controller. The triggers on the Touch controller are force-sensitive, able to detect hundreds of discrete positions.

From these data, the Oculus software can reverse engineer a virtual reality avatar of the user's hand, a "hand model." The hand model independently articulates the thumb and index finger from the other three fingers (middle, ring, and pinky). The controllers have tracked optically through the "Constellation" a series of infrared (IR) light-emitting diodes (LEDs) that are detected and tracked by cameras either mounted on the wall (outside-in) or the headset (inside-out).

For the original Oculus Touch, each controller contains a single Alps Haptic Reactor two-dimensional linear resonant actuator (LRA). The Alps Haptic Reactor contains two resonant frequencies -- 160 in the horizontal (x-axis) and 320 in the vertical (y-axis). The linear oscillator is mounted low down in the grip, near the battery, as it would deliver the most effective vibration there, while also helping to achieve the Touch's excellent weight distribution. By allowing two axes of freedom, the Haptic Reactor can approximate fully realistic waveforms. The

> “TO MAKE TRULY TANGIBLE
> IVES, THE TOUCH CONTROLLER
> LINE-UP NEEDS AN UPGRADE.

frequency response of the Haptic reactor does, however, limit certain interactions. Unless the interaction occurs at either of the two peak waveforms, the haptic signal loses fidelity -- weakening the experience for the user.

The newest Oculus Touch Controller -- the “Quest Touch” or LCON controller launched with the Oculus Rift-S and Oculus Quest HMDs. A simplified design of the original Touch controller, the Quest Touch controller relies instead on a smaller single LRA sacrificing even more haptic fidelity for lower production cost and ease of manufacturing. The Quest Touch also introduces inside-out tracking, with a ring of IR LEDs -- the Oculus Constellation -- being repositions from down around the knuckles, to an upward position that allows the HMD mounted cameras to track the controller.

For a successful IVE, the technical elements of the hardware and software must come together to convince your brain that what you’re virtually experiencing is actually real. Without immersion, the illusion will feel incomplete. While the Rift HMD can teleport you into another world, if your hands don’t make the trip, the experience will fall short. Both the Touch and Quest touch controller help bring our hand into the IVE, but there are many interactions that are simply not possible with the hardware limitations of these designs. To make truly tangible IVEs, the Touch controller line-up needs an upgrade.

Fig: Oculus Quest Touch “LCON” Controller

THE CHALLENGE, THE SCOPE, & THE REQUIREMENTS

CHAPTER 02

THE CHALLENGE, THE SCOPE AND THE REQUIREMENTS

THE CHALLENGE: FOUR DESIGN FACTORS TO PRODUCE A SUCCESSFUL HAPTIC PRODUCT

The task of producing such an upgrade would not be easy. To establish some sort of framework, around the design space, I wanted to use some simple design factors that would help me evaluate ideas. I'd developed the idea of a balancing act between three factors -- performance, comfort, and realism -- when developing another haptic interface for virtual reality, the SenseGlove[1]. To these, I added in a fourth -- immersion -- something I'd previously blended into realism. These four design factors would give me a sense of the trade-offs needed to build a product for an IVE.

The first design factor -- performance -- is the durability, stability, and repeatability of an interaction presented by the controller. A high-performance controller can be used without issue and repeated re-calibration. A high-performance interface can portray the desired interaction to a large range users in various positions, grips, and orientations. It should work the same over hundreds of cycles. In short, a high-performance interface is predictable and repeatable.

The second design factor -- comfort -- is reflected in the user's ability and desire to use the Oculus touch controller for extended periods of time. A comfortable controller fits small hands and large hands without issue. A comfortable controller is padded and non-disruptive. A comfortable controller is lightweight and easy to maneuver. It fits seamlessly into the mechanical interaction, not noticeably altering the interaction style adopted in everyday interactions. Finally, a comfortable interface produces no discomfort. In short, a comfortable controller is a joy to hold, wear, and/or use.

The third design factor -- realism --is driven by the user's ability to experience the feedback as if it were the predicted outcome of the interaction. A realistic controller can portray experiences in IVEs in a precise and accurate manner. Button clicks produced by a realistic controller mirror those of the real world. A digital grab should feel like a physical one. Haptic feedback that is not realistic interferes rather than enhances the experience. In short, a realistic controller makes an IVE feel real.

The fourth design factor -- immersion -- is closely related to, but distinct from realism. As part of my reflection of the SenseGlove project I'd been working on since the AED course at TU Delft, I'd come to realize that in IVEs realism was different than

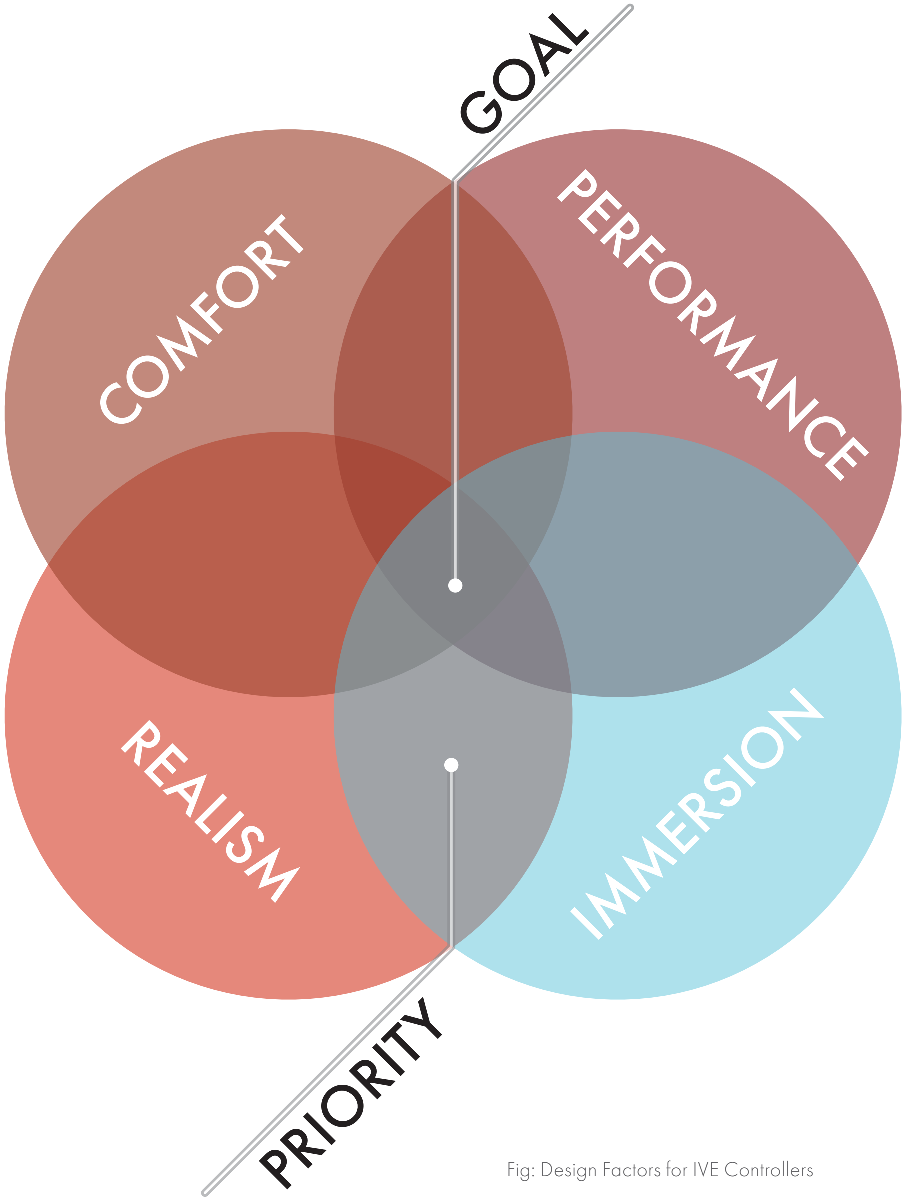

Fig: Design Factors for IVE Controllers

immersion. Realism was the accuracy of the experience, while immersion was the believability. The central tenet of immersion is reducing the "noise" of the controller and producing the clearest haptic feedback signal to maximize believability. An immersive controller should allow the user to adapt -- The product should fall away after a few minutes of use and not be perceived. An immersive controller augments the cues the IVE presents for a broad set of user scenarios. If an IVE is an illusion presented to the user, an immersive controller makes the user forget they're using anything to interact with that illusion.

THE SCOPE: PRIORITIZING AN INTERACTION

Rather than try to tackle all interactions in the world, a task far greater than this thesis would allow, I further narrowed the design space of my thesis by proposing a list of twelve interactions that felt relevant and useful in IVE.

The first category was controls. The interactions in this category all represented a virtual recreation of a real-world control -- buttons, keyboards, or knobs. Comprising of over half of the interactions proposed, these six interactions aimed to help users of an IVE better understand the digital replications of these real-world objects. The primary use case for this was virtual training, a common use for VR technologies that allowed users to learn how to control complex machines and systems with little risk. The interactions were as follows:

- Button Press (clicky)
- Button Press (soft)
- Keyboard Press
- Knob Turn (Stepped/Castellated)
- Knob Turn (Smooth/Continuous)
- Slider Move / Linear Potentiometer

The next category was two interactions that presented the first consumer use case for IVEs, gaming. I looked at what sort of games and interactions were common in popular VR video games SuperHot VR and StarTrek Bridge Crew. These two games represented two popular categories of game, the first-person shooter (FPS), and the Simulation genre. For FPS's, the primary interaction was unsurprising, shooting, and

for the Simulation genre, there were lots of futuristic touchscreen and holographic screen displays. Thus for the gaming category, the two following interactions were chosen:

- Trigger Pulls
- Touchscreen Press

Third were assembly tasks. Useful for both of the prior use cases, these interactions represented two virtual objects interacting with each other. Our sense of touch is essential in assemblies and tool based interactions, transmitting force, and torque information to our fingers and hands. I wanted to be able to know when an object I was holding hit the world, or when The ability to make tools feel right was useful in both gaming and virtual training. To cover these interactions, I chose the following:

- 2nd Order Effects (Tool to world interactions)
- Insertions / Snap Fits

The fourth and final category was motions. One of the biggest issues in VR is that digital objects don't carry the same physics properties as their real world counterparts. A digital object has no inertia, no mass, and no friction. All of these physical properties are a key part of the experience in manipulating objects. We know how hard to grip an object from the stick slip in our hand. Without these properties, or cues of these properties, virtual objects don't feel real. To help resolve this issue, we proposed two motions

- Stick/Slip in Hand
- Weight/Inertia

THE SCOPE: PRIORITIZING AN INTERACTION (CONT'D)

With twelve potential interactions, I had to downselect even further, choosing an interaction that was sufficiently challenging to accomplish in VR while still being possible within the time limit of the thesis. I also wanted an interaction that really could be improved, something that felt limited by current technologies. One interaction that exemplifies this limitation is that of button interaction.

Buttons have been a core part of human-machine interfaces for centuries. Indeed, even as we have moved from analog to digital through the computer age, buttons -- the mouse and keyboard -- have come to define how we interact with electronic devices. Newer interface paradigms such as touchscreens still emulate the haptic interaction of traditional button profiles.

As an example, when interacting with an iPhone 7 "home button", the "button" is a flat glass plate, the "click" haptic queue is generated by the LRA inside the Taptic Engine. Even though there's no displacement during the interaction, the prevalence of buttons demands the emulation of their physical properties.

With an interaction set, I needed to determine the other constraints on my project scope.

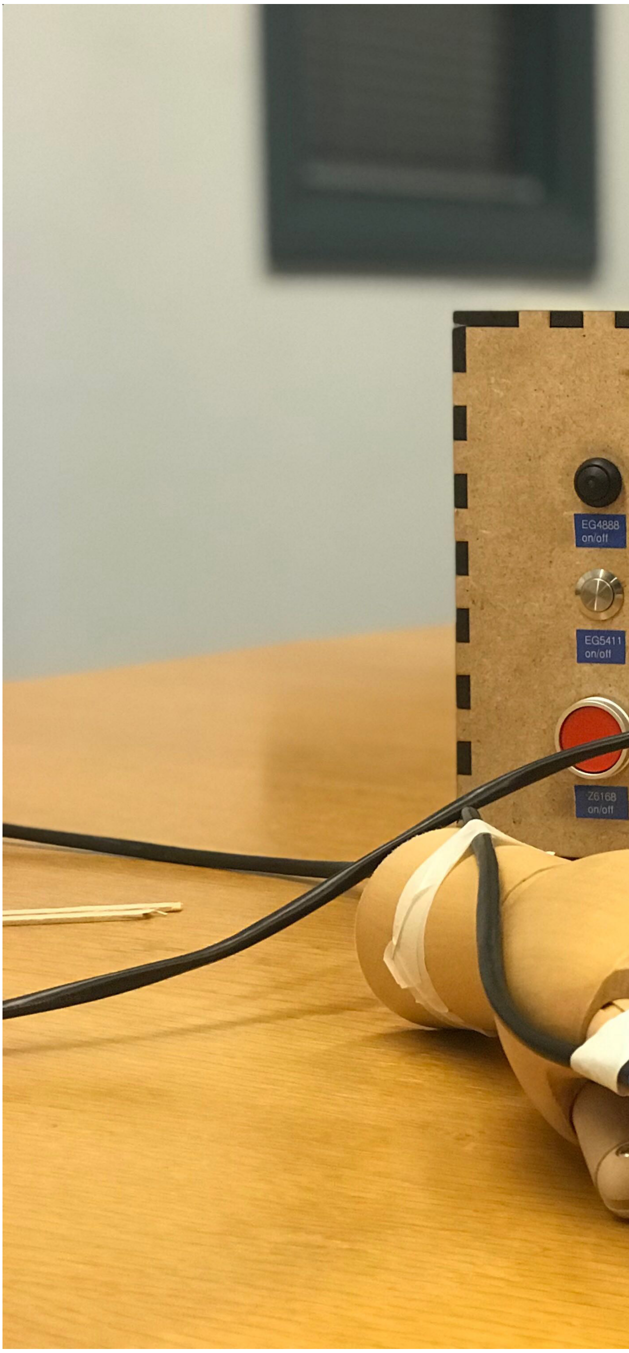

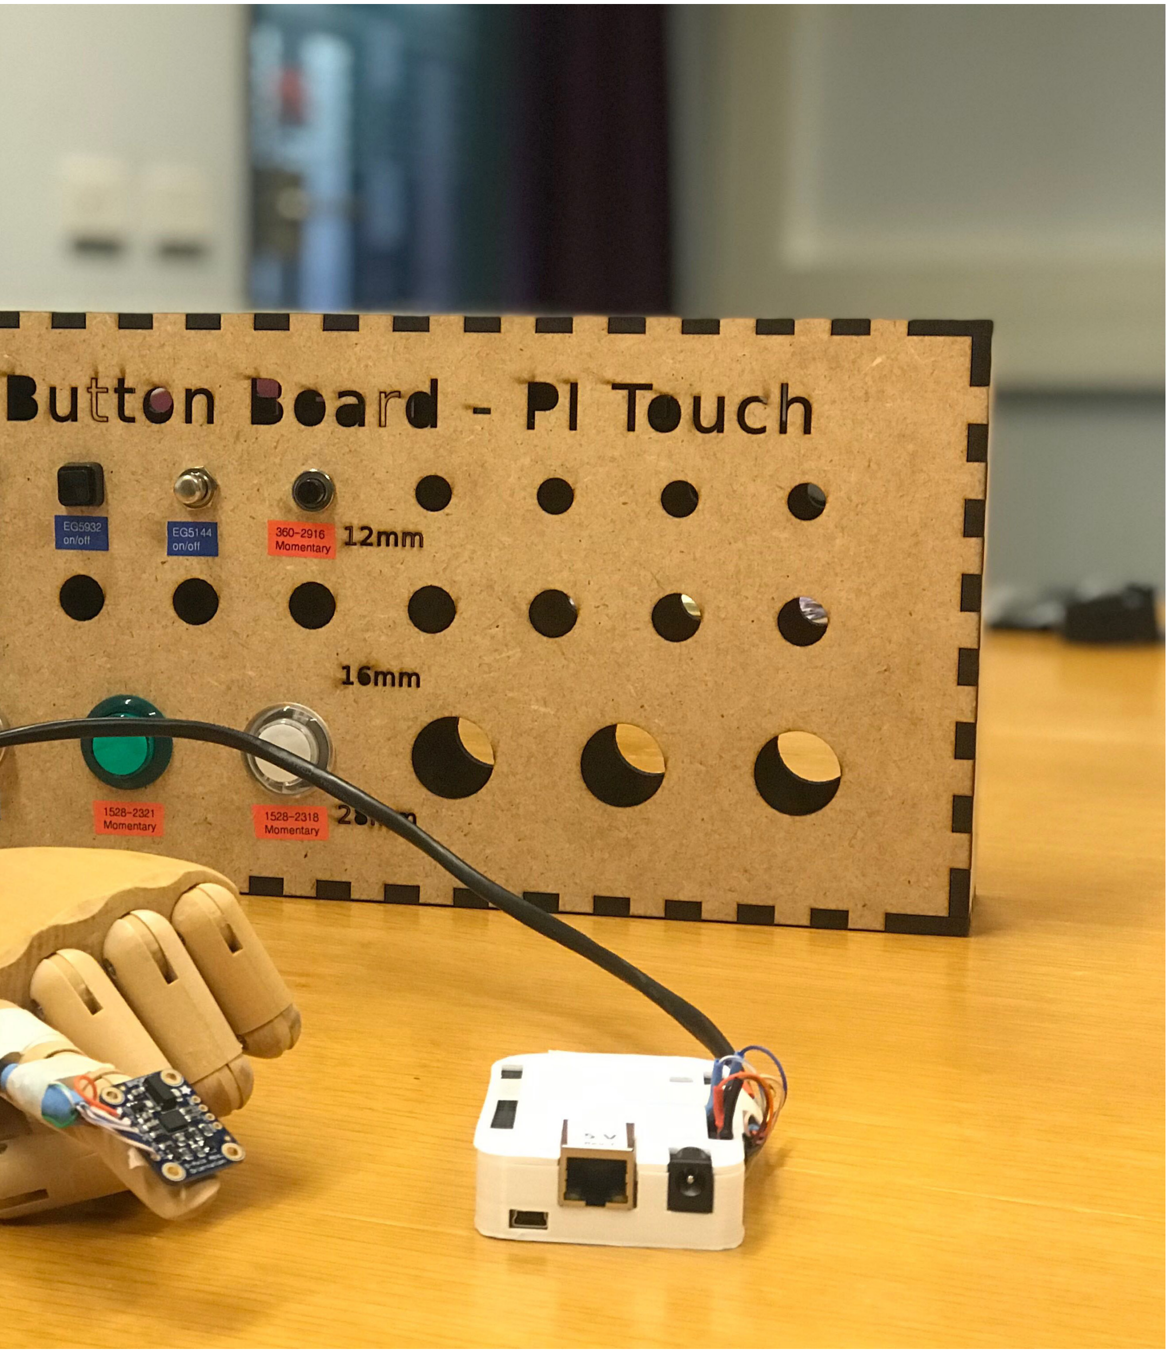

Fig: Button Box, containing various styles of tactile button for actuation recordings. ADXL 345 Haptic Recorder Mounted to Hand

THE REQUIREMENTS: LIMITING THE DESIGN SPACE

Like many creatives, I work best with a strong set of limits. It keeps my mind from wandering, and helps me stay focused on the core of the design. It's easier to make decisions in a defined solution space. This would be my last tool to help contain the design space. Through user interviews, expert interviews, and my own experiences, I would come up with a comprehensive list of requirements.

Great products are built from great plans. Defined and documented requirements are a key part of the process for the development of a new or complex system. To ensure the product meets users' needs, it needs to be understood, captured, and agreed upon. Good requirements can help us see around corners -- we can find any intrinsic conflicts in the design -- and can help the roadmap of the product design process.

To do this, I built two requirements documents, the first, an engineering Product Requirements Document [2], which helped explore the regulatory, serviceability and electronic aspects of the product. The second, a more traditional weighted list of requirements follows closer to the Delft Design Guide outline and helped me manage stakeholder and contextual requirements. [3]

KEY REQUIREMENTS

The first requirements I came up with revolved around the desired interactions. I chose four that I felt were an absolute "must" given the discussions with the leadership of this project on the Oculus side, and the use cases we'd defined. These requirements would help solidify the realism and immersion design factors.

› The New Controller must be able to accurately render a "clicky" button press
› The New Controller must be able to accurately render a gun recoil
› The New Controller must be able to accurately render a trigger pull
› The New Controller must be able to accurately render footsteps

The next set of requirements case from my experiences with haptics. From my prior projects, and research outlined in section "The Biology of Haptics", I knew that haptics was experienced in a range from 20-1kHz. However, what I called the haptic "core" frequencies ranged from 60hz to 400hz. Higher frequencies would be barely perceptible, but could be used to produce some novel effects. These requirements reflect the three basic generic waveforms across what I felt were the most useful ranges of frequencies in haptics. These requirements would define the performance design factor.

› The New Controller must be able to

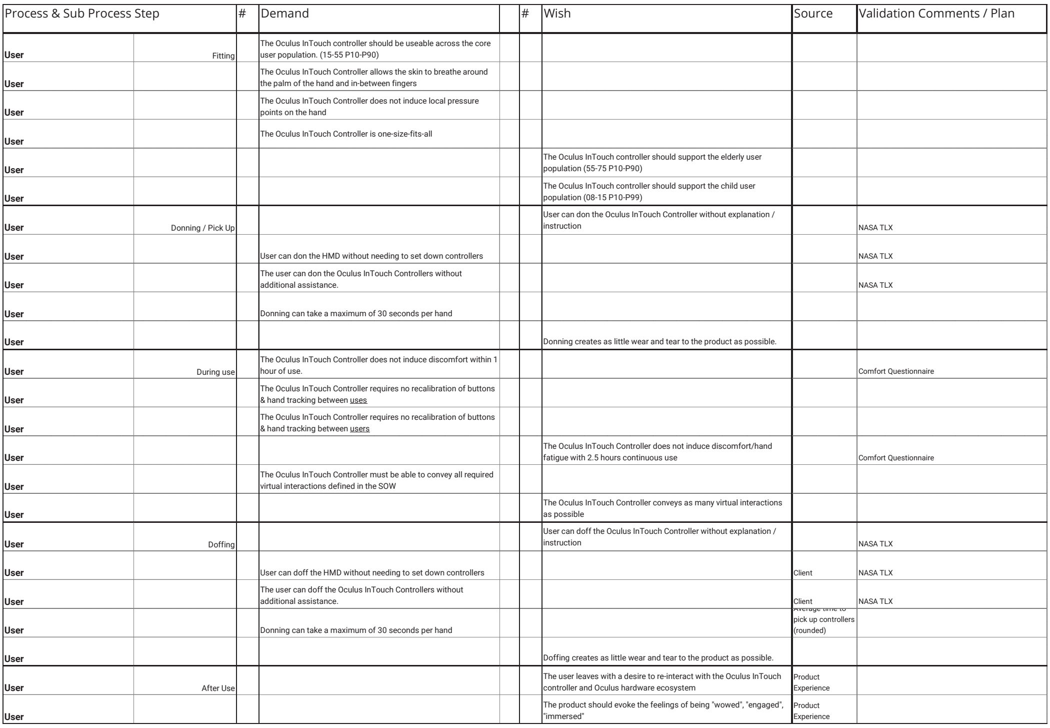

Process & Sub Process Step		#	Demand		#	Wish	Source	Validation Comments / Plan
User	Fitting		The Oculus InTouch controller should be useable across the core user population. (15-55 P10-P90)					
User			The Oculus InTouch Controller allows the skin to breathe around the palm of the hand and in-between fingers					
User			The Oculus InTouch Controller does not induce local pressure points on the hand					
User			The Oculus InTouch Controller is one-size-fits-all					
User						The Oculus InTouch controller should support the elderly user population (55-75 P10-P90)		
User						The Oculus InTouch controller should support the child user population (08-15 P10-P99)		
User	Donning / Pick Up					User can don the Oculus InTouch Controller without explanation / instruction		NASA TLX
User			User can don the HMD without needing to set down controllers					NASA TLX
User			The user can don the Oculus InTouch Controllers without additional assistance.					NASA TLX
User			Donning can take a maximum of 30 seconds per hand					
User						Donning creates as little wear and tear to the product as possible.		
User	During use		The Oculus InTouch Controller does not induce discomfort within 1 hour of use.					Comfort Questionnaire
User			The Oculus InTouch Controller requires no recalibration of buttons & hand tracking between uses					
User			The Oculus InTouch Controller requires no recalibration of buttons & hand tracking between users					
User						The Oculus InTouch Controller does not induce discomfort/hand fatigue with 2.5 hours continuous use		Comfort Questionnaire
User			The Oculus InTouch Controller must be able to convey all required virtual interactions defined in the SOW					
User						The Oculus InTouch Controller conveys as many virtual interactions as possible		
User	Doffing					User can doff the Oculus InTouch Controller without explanation / instruction		NASA TLX
User			User can doff the HMD without needing to set down controllers				Client	NASA TLX
User			The user can doff the Oculus InTouch Controllers without additional assistance.				Client	NASA TLX
User			Donning can take a maximum of 30 seconds per hand				Average time to pick up controllers (rounded)	
User						Doffing creates as little wear and tear to the product as possible.		
User	After Use					The user leaves with a desire to re-interact with the Oculus InTouch controller and Oculus hardware ecosystem	Product Experience	
User						The product should evoke the feelings of being "wowed", "engaged", "immersed"	Product Experience	

Fig: Screenshot from List of Requrements for Oculus inTouch Controller

render a sine-wave from 60hz-600hz

› The New Controller must be able to render a sawtooth-wave from 60hz-600hz

› The New Controller must be able to render a square-wave from 60hz-600hz

› The New Controller must be able to repeat across 1 million cycles any of these waveforms with no significant deviation in amplitude or frequency

I had a few requirements to support comfort. A controller that was too heavy, slippery or poorly balanced would be a slog to use. The goal was to create an improved controller, but those improvements couldn't impact the usability.

› The New controller must not weigh more than 160 Grams (w. Battery)

› The New Controller must have a center of gravity located between the user's middle and index fingers.

The final set of key requirements related to the user experience. I had a vision for how the user should experience the produce. The user experience goals and market position goals would help constrain the design from becoming too much of a R&D project and keep me firmly grounded in what was possible and implementable.

› The product should evoke the feelings of being "accessible", should be "flexible across use cases" and provide a "compelling experience"

› The MSRP for the New Controller must remain at $75 / hand or $150 /set

MY SUPPORT TEAM: EXPERTS & ADVISORS

Fortunately, I wasn't alone in trying to solve the challenge of the Oculus inTouch controller redesign. I had a team of instructors, industry experts, researchers, and fellow students to support me:

JESS HARTCHER-O'BRIEN

CHAIR
HARDWARE DESIGN AND ENGINEERING

HENK CRONE

MENTOR
PRODUCT DESIGN ENGINEERING

ROBIN MILLER

COMPANY MENTOR
HUMAN PERCEPTION, HAPTICS, AND MULTI-SENSORY

HENRIC JENTZ

COMPANY MENTOR
HARDWARE DESIGN, INDUSTRIAL DESIGN

MARK ANDO

PERSONAL MENTOR
HARDWARE DESIGN AND ENGINEERING

MICHAEL WIERTELESTKI

UNIVERSITY EXPERT
HAPTICS AND ACTUATOR ENGINEERING

LINDA PLAUDE

UNIVERSITY EXPERT
SMART TEXTILES

MARTIN HAVRANEK

UNIVERSITY EXPERT
PCB DESIGN

BEN SCHNIDERS

LOFELT GMBH
DSP DESIGN & LOFELT PRODUCT

RAFAL PIJEWSKI

ACTRONIKA SAS
ACTUATOR ENGINEERING & ACTRONIKA PRODUCTS

KARINA DRILLER

PITOUCH LAB
NEUROSCIENCE OF HAPTICS

BRYAN ZAAIJER

PITOUCH LAB
PRODUCT FORM & CMF

MICHELLE CORTEN

PITOUCH LAB
HAPTICS RESEARCH & DEVELOPMENT

SIMON IMMERZEEL

STUDENT
CONCEPT ART & SKETCHING

A PREVIEW: THE OCULUS INTOUCH CONTROLLER

CHAPTER 03

THE OCULUS INTOUCH CONTROLLER

The goal of this project was to take a magic trick -- the illusion of an IVE, and make it more compelling through improved touch. Months of research, experiments, and dozens of quick prototypes or haptic sketches have come together to produce a product I call the Oculus inTouch controller.

After one hundred days of hard work a complex prototype is finished and functional. The inTouch controller navigates the balance between the key design considerations -- comfort, performance, realism and immersion, and improves on the Oculus Touch and Quest Touch controllers. The inTouch concept lives up to its name, helping users stay in touch with the digital world and each-other through new and improved haptic feedback.

Fig: Oculus InTouch renders.
Top: Exploded view
Bottom: Side view.

RESEARCH: THE BODY-CENTRIC APPROACH TO HAPTICS

CHAPTER 04

RESEARCH: THE BODY-CENTRIC APPROACH TO HAPTICS

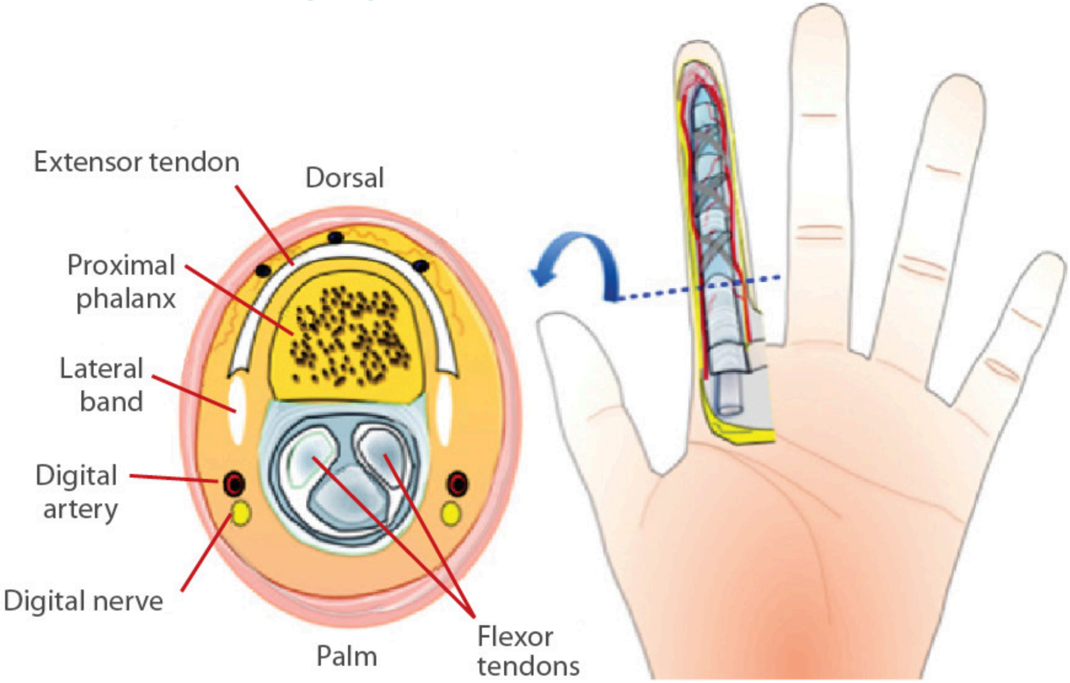

Fig: Finger Cross Section

WHERE TO BEGIN?

To date, haptics has come from two primary schools of thought -- the body-centric approach, and the technology-centric approach. The body-centric approach isolates the body while presenting a series of stimuli. By exploring the responses, the body-centric approach aims to understand the sensory I/O of the human system. The technology-centric approach ignores the user and instead attempts to engineer hardware that will reproduce a particular waveform or pattern.

Both of these approaches have long had a place in haptics. To design the Oculus inTouch controller, I began by exploring both. I wanted to understand the limitations and weaknesses to each of these schools of thought, and to see which of the two would better lead me to success. One, the body-centric, played to the background of my support team at the TU. The other, the tech-centric, appealed to my years of experience as a mechanical engineering consultant in the US.

The drive to understand touch has typically been driven by the human desire to understand better what sits below our skin. Anatomy and biology are the foundation

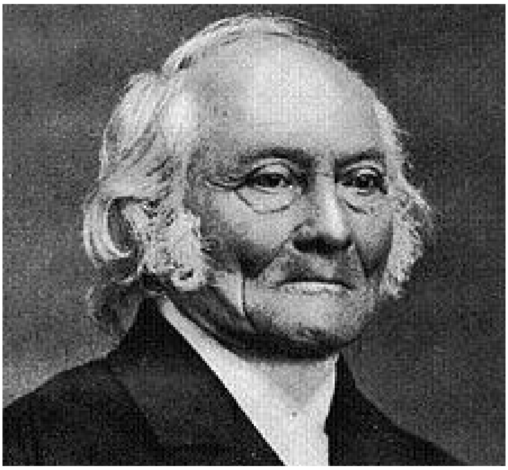

Fig: Weber (Left) & Fechner (Right)-- Grandfathers to Psychophysics

for our exploration of touch, with Weber's [4] and Fechner's [5] mid 19th-century experiments providing a foundation for what is now known as psychophysics.

Psychophysics -- the study of the relationship between physical stimuli and sensations/perceptions induced defines the body-centric approach. Pulling from both the biological and psychological field, experimental psychophysics explores sensory perception from an isolated, and controlled perspective.

Gescheider's early explorations on static touch , helped identify the wildly complex anatomical pathways between our hands and our brain. Dozens of papers outline the linkages between mechanoreceptors and brain [6]–[13]. Our sense of touch is derived from both cutaneous and kinesthetic inputs, where both passive and active interactions generate changes in the mechanical state of our bodies and the world.

> *"THE BODY-CENTRIC APPROACH IS FOCUSED ON ANSWERING THE "HOW" OF HAPTICS -- HOW DOES OUR BODY UNDERSTAND AND INTERPRET THE HAPTIC INPUTS.*

The body-centric approach is focused on answering the "How" of haptics -- how does our body understand and interpret the haptic inputs. The isolated, controlled approach that psychophysics offers is excellent for understanding in a research environment the input and output (I/O) of both brain and body.

However, our experience of the mechanical world depends on an ability to move and manipulate. As such our perception derived from the sense of touch depends on the action/interaction phase of a movement, the goal of the observer, among other factors. [11]

BIOLOGY OF HAPTICS

Our brain relies on each of our five senses to provide a coherent and robust perception of reality. Each sense functions by recognizing patterns of stimulation on unique organs -- our eyes detect light waves, our ears sound pressure waves. Our sense of touch or "haptics" comes from a variety of sensors spatially distributed throughout the body and interwoven among our motor mechanisms and functions. Touch is primarily driven through our hands, with the fingertips and palms among the most sensitive areas in the human body[14]. Humans successfully use their hands to identify objects and to extract information -- such as surface texture, compliance, weight, shape, size, orientation, and thermal properties[14], [15]. Current haptic interfaces provide two primary forms of haptic information.

First, there is kinesthetic feedback -- information about the position and motion of the hand relative to an object. These sensors are tied deeply into the motor system of the human body. Muscles, tendons, and joint sensory receptors are stimulated or affected to produce kinesthetic feedback. Kinesthetic feedback is useful in human-computer interfaces to provide information about an object's shape, size, and mechanical properties -- stiffness, strength,and elasticity.

> *OUR SENSE OF TOUCH OR "HAPTICS" COMES FROM A VARIETY OF SENSORS SPATIALLY DISTRIBUTED THROUGHOUT THE BODY AND INTERWOVEN AMONG OUR MOTOR MECHANISMS AND FUNCTIONS.*

Second is tactile feedback, providing information on the geometry, texture, friction of an object. To produce tactile feedback, tactile sensory receptors known as mechanoreceptors are stimulated [16] [17]. These receptors detect only the skin deformation or vibration caused by

	Properties Sensed	Location Within the Hand	Frequency Response (Hz)	Response / Adaptation	Receptive Field / Range	Depth (Type)	Shape
Merkel's Disks	Shapes, Pressure, & Edges	Hair Follicles, Fingerprint Ridges	5 - 15	Slow (Sustained)	Small	Shallow (Type I)	Disk / Dome
Ruffini Endings	Skin Stretch & Elongation	Joints, Fingerpad	15 – 400	Slow (Sustained)	Large	Deep (Type II)	Long Spindles
Meissners Corpuscles	Light Touch	Thin, Hairless Skin (Finger Pads)	20 – 50	Fast (Pulsed)	Small	Shallow (Type I)	Conical / Domed
Pachinian Corpuscles	Vibrations	Hairless, fatty tissue.	60 - 400	Fast (Pulsed)	Large	Deep (Type II)	Ovaled / Cylinder

Table: Detail on mechanoreceptor properties.

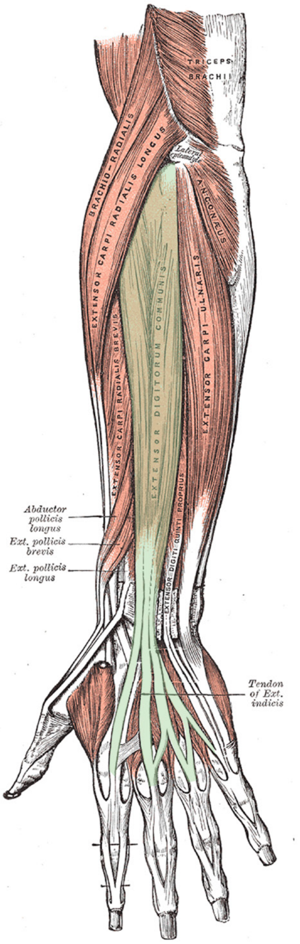

Fig: Musculature of the hand and forearm

contacting objects [18].

There are four types of tactile receptor: Merkel discs, Meissner corpuscles, Pacinian corpuscles, and Ruffini endings. Merkel discs and Meissner's corpuscles lie at the epidermis/dermis interface. They resolve textural features of 0.6-1mm in spacing. Merkel discs are activated in static interaction, and respond to normal loads and pressure, resolving the surface area of interaction. Meissner corpuscles are activated in dynamic interactions. They enable feature edges to be detected during sliding contact. Pacinian corpuscles and Ruffini endings lie closer to the core of the finger. Pacinian corpuscles cannot resolve locations accurately but detect higher frequency (200-300hz) vibrations than Merkel discs or Meissner corpuscles[19]. Ruffini endings detect tangential forces and stretch in the skin[20], [21] -- see below right for an outline of the various mechanoreceptors and their properties.

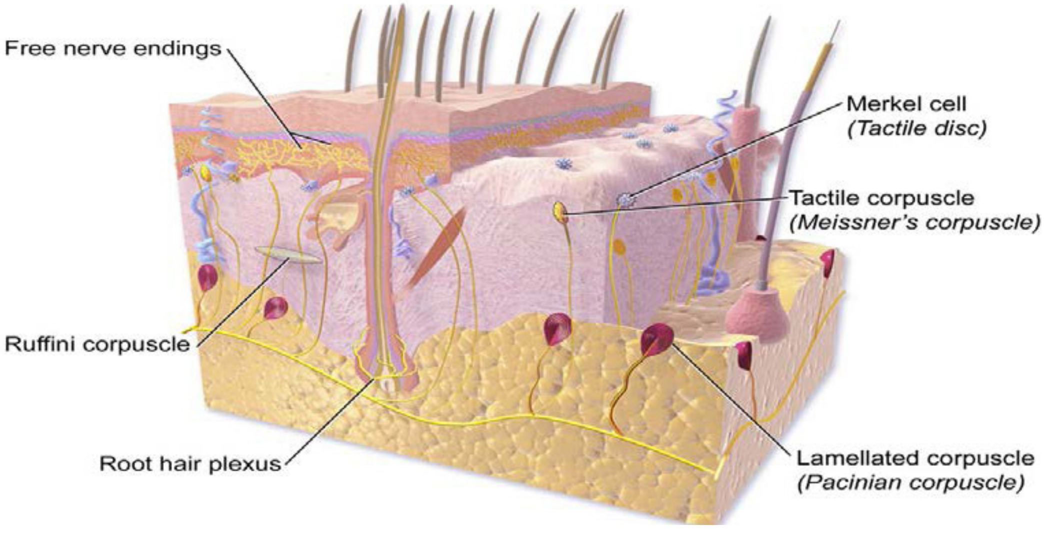

Fig: Skin cut-away showing the location and shape of mechanoreceptors inside of the skin.

PERCEPTUAL LIMITS OF TOUCH

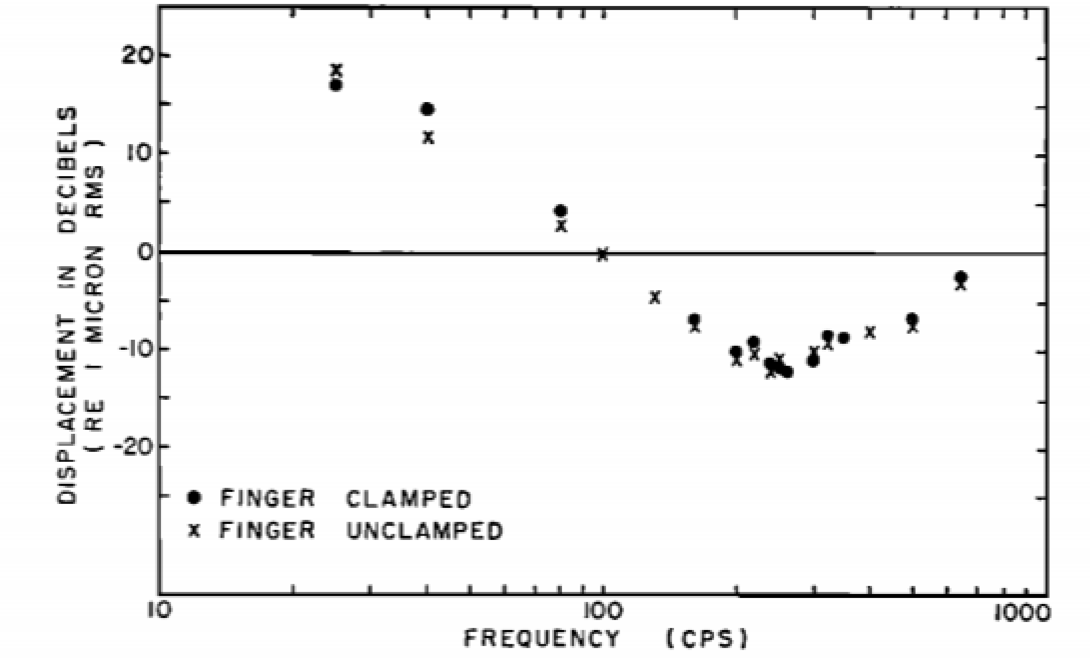

Fig: Frequency response of hairy skin, reprinted from Verillo 1962.

The obvious limit to touch is the range -- we have to be in contact with an object in order to explore it with touch. Thus, anything outside the human reach envelope (just over a meter horizontally) is outside the perceptual range.[22]

Even for objects within reach, there are limits -- as only areas of contact can convey information. As opposed to vision, where an object is always viewed with the environmental context, haptic stimuli must be scanned to perceive context. This scanning requires a great deal of cognitive effort, and comes with a substantial error value -- try drawing out a map that you've only ever felt with touch! Touch is a relative sense -- it detects change -- and not an absolute sense -- one with a fixed origin. When we look at the work through our eyes, our eyes behave as a reference point, an origin. This changes when we explore touch, with various parts of the body behaving as the reference.[23]

Touch is the first sense to develop in humans, and as children, touch is our dominant method of exploring the world. We rely on touch, as it can provide a ground truth is speedy and reactive.

Neuroscience has established that the haptic sense can discriminate two stimuli within 5ms [9], [10]. Touch reacts faster than our sense of sight, triggering reflex actions before conscious thought. Fortunately, the granularity of touch is lower than the response threshold. Touch sensations can lag as much as 250ms from other stimuli and still be causally linked, though the true

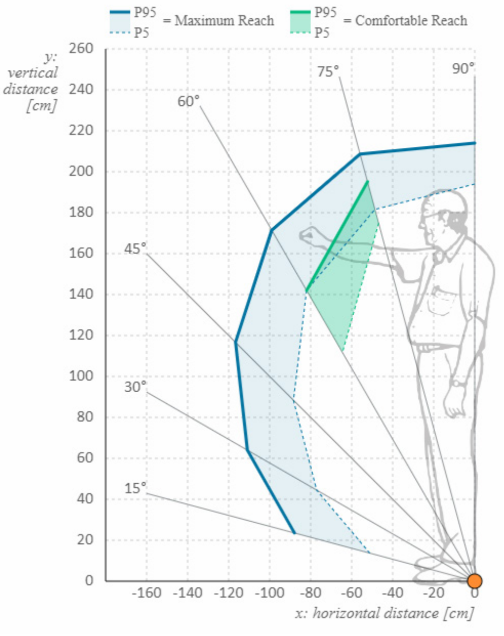

Fig: Male reach envelope, from DINED

Fig: Perceptual temporal limit between visual & haptic cues

working range is below 75ms.

Our sense of touch is also capable of fine spatial discrimination. Haptically we can discern a displacement as small as 0.2 microns in length on the fingertips[24]. We can detect a height change of 0.1 microns with the pads of our fingertips.

Touch is a low-frequency sense. We can perceive frequencies down into the single digit range, and above 1kHz. However, as shown by Verillo[25], [26], mechanoreceptors peak in sensitivity between 200-300Hz, with the useful working range somewhere between 60-400Hz.

Interestingly, touch is biological-sex dependent. Consistently men are shown to have worse perceptual thresholds than women in both spatial[27], and frequency domains[28]. It is also age-dependent, with touch sensitivity decreasing as we age and our bodies deteriorate. This makes design for haptics challenging.

We typically design a product around the lowest common denominator -- e.g. a controller must fit the smallest hand in our user group. However, with VR, what is designed around the core demographic -- young men -- may feel oversaturated to young women. This exact same sensation can feel weak to the elderly.

THE COMPLEXITY OF TOUCH

AN INTEGRATIVE SENSE

To the layman, touch is often viewed as a single comprehensive sense. In truth, our sense of touch is several modalities/ perceptual channels being integrated by the brain into a single picture. Touch is traditionally broken down into thermoreceptive, nociceptive, kinesthetic, and mechanical sensing channels [29]. Each channel plays a crucial role in producing a holistic sense. The figure across illustrates some of the perceptive channels stimulated in the interaction with a hole punch.

Thermoreception is our sense of hot, and our sense of cold. Expectations drive thermoreception -- we expect a shaded stone to be cold, and metal near a fire to be hot. We pull from our memories and other sensory cues to determine our expectation of temperature. Thermal sensing allows us to discriminate and identify material properties of objects, from which we can tune manipulation behaviors accordingly. Additionally, of all our haptic senses, thermoreceptors have some of the strongest hysteresis -- they remember their former state.

Nociception relates to the encoding and processing of damage. It relies on nociceptors (pain receptors) that can detect mechanical, thermal, or chemical damage through pain. Nociceptive signals drive our reflex reactions and are a crucial part of the haptic system primary purpose -- to protect the body. In IVEs, Nociception drives our natural bias towards rejection of the illusion.

Kinaesthetics describes the perception of one's body motions through the sensors within our muscles and joints. Based on proprioception— one's own conscious and unconscious perception of the forces, torques, movements, relative positions and angles of neighboring parts of the body, kinaesthetics helps us understand the weight and shape of objects.

Tactition, is the perception of direct contact and relative motion between the skin and the objects of interest. Tactition can further be broken down into cutaneous perception -- tactile information from skin pressure and skin deformation -- and vibrotactile perception-- information from vibration waves propagating through the skin. Through these channels, we understand shape and texture.

When we interact, we take perceptions of each of these different sensory channels and combine them to produce a cohesive tactile picture of the interaction. In the body-centric approach to haptics, the idea is to isolate each channel and to track the effects of a stimulus on a single channel. When we need to design a product that utilizes touch, we, however, need to ensure the opposite, designing for the holistic experience, targeting the correct sensation to each of channels to produce a clean, clear, comprehensive experience.

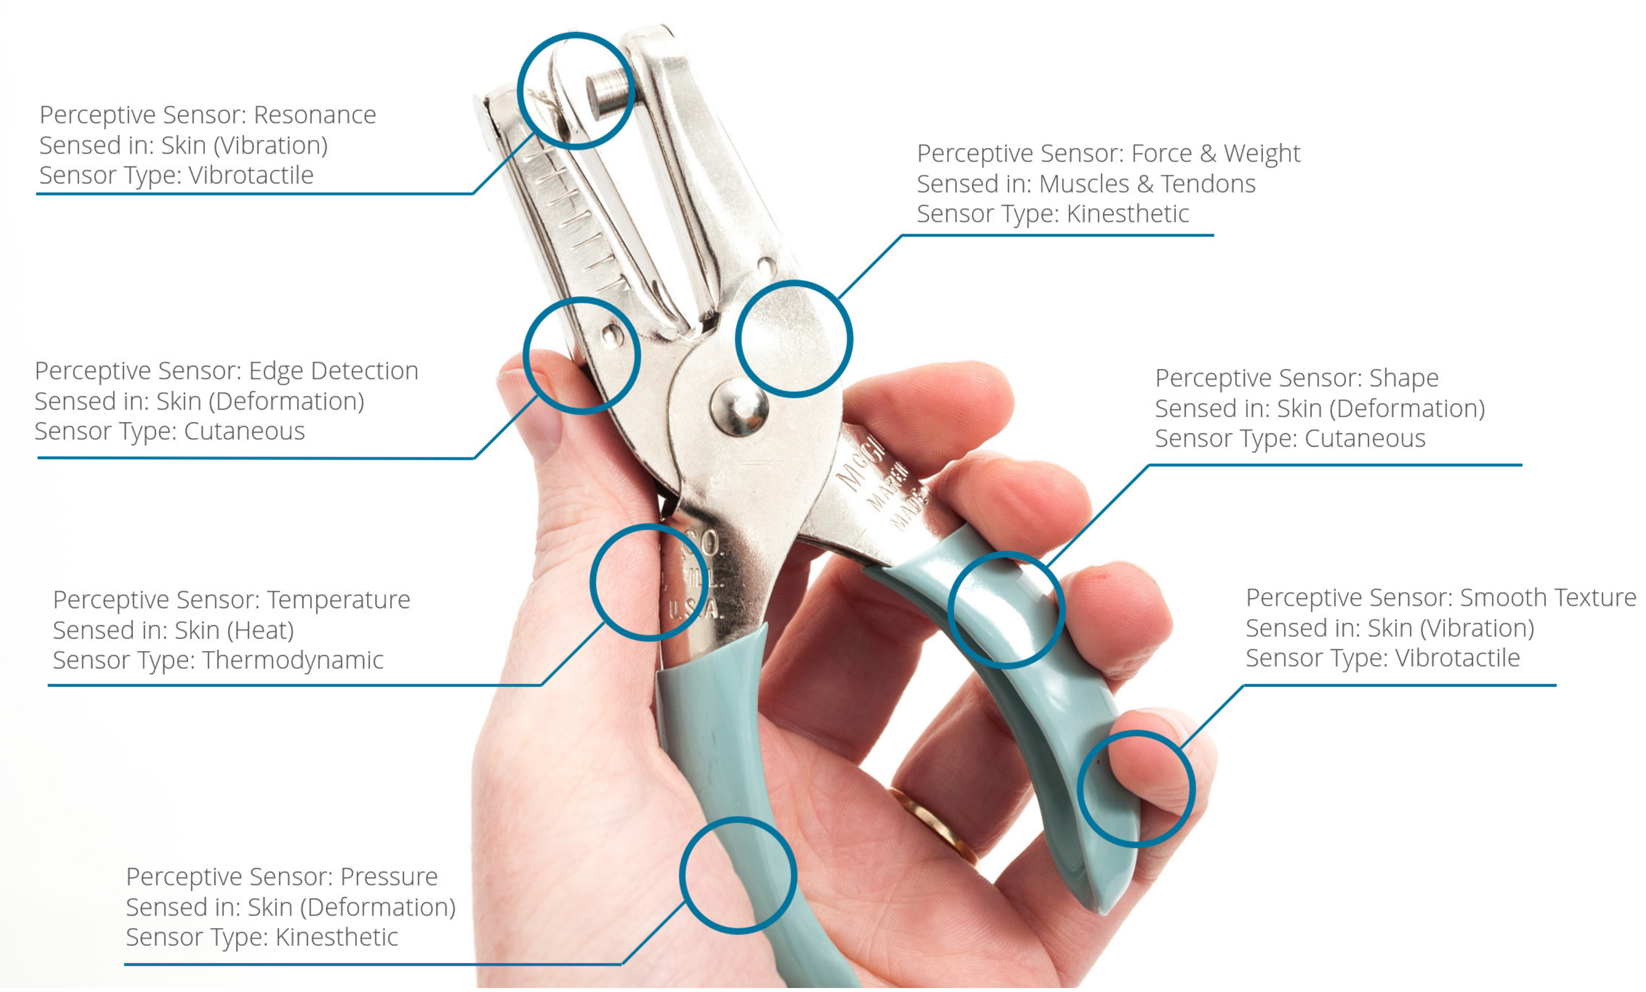

Fig: Different aspects of touch as shown through a hole punch interaction.

THE COMPLEXITY OF TOUCH: AN ACTIVE SENSE

Our brain is active in touch – interpreting all these perceptual channels into a single cohesive picture. While it is possible to create immersive and even realistic touch simulacra without complete engagement, each missing channel makes those simulacra feel hollow or somehow "wrong." There is room for interpretation, and therefore, room for mistakes.

There are dozens of "haptic illusions" that show the ambiguity in our sense of touch. These illusions highlight possibilities to trick the brain/user into perceiving/having an experience without providing the exact physical inputs that typically drive the experience. From Hayward, "Take an ordinary comb and pencil and lay your index finger along the top of the comb, then run the pencil back and forth along the side of the teeth. Even though the teeth are moving from side to side in a wave-like motion, your finger will feel as if a raised dot is traveling up and down the comb.

A similar effect happens if you take two coins and put them in the freezer until they are thoroughly chilled. When this is done, take them out of the freezer and place them on a table - one on either side of a coin that's room temperature. If you put your ring, middle and index fingers on the coins, even though the ring and index fingers are the only ones touching a cold surface, all three fingers will feel chilled.

> *IF SOMETHING SHOULD FEEL A CERTAIN WAY, OUR BRAIN WILL TAKE LEAPS TO MAKE IT SO.*

In touch, our brain is trying to fill in the gaps in our sensory picture, and will often take a guess. If something should feel a certain way, our brain will take leaps to make it so. To do this, the brain relies heavily on other sensory data -- vision, sound, etc. -- but also on memory. It is why we recoil at a strangers touch, but the same touch

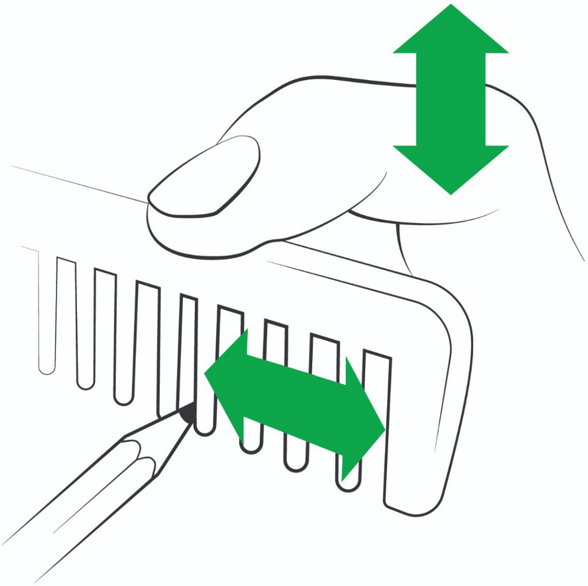

Fig: Hayward's haptic illusion with a comb.

stimuli from a friend or lover can comfort. It is also why we release rotten fruit before we cognitively realize it has rotted. Our brain will fill in gaps to build what should happen-- the expected outcome -- but is hyper-aware for any miscues or noise in the signal. Therefore, our cultural background, upbringing, and unique set of memories make touch unique for every person.

> *IF WE DE-COUPLED REALISM FROM IMMERSION, AND GAVE UP ENTIRELY ON THE IDEA OF REPRODUCING THE EXISTING WORLD, WE COULD DO MORE WITH FEWER ACTUATORS.*

This 'best guess' process of perception provides potential for exploitation in recreating the world of touch. A clever IVE takes advantage of this, presenting only the minimum amount of cues necessary in the cleanest, clearest manner possible. Rather than try and recreate every sensory detail for peak realism, presenting only the most essential stimuli and allowing the brain to fill in the blanks can create higher immersion. This approach, which I call haptic cuing, takes advantage of the natural biases of the brain to allow the IVE to self-improve. The more time spent in the environment, the more accurate the brain's prediction of the environments behaviors, and thus, the more it believes the illusion.

For me, this knowledge was one of the most useful things I learned in the exploration of the body centric approach. The idea that haptics had an information hierarchy wasn't new, but Jess and I had yet to see work that truly leveraged this idea and applied it to a VR/IVE context. Biology had revealed to us a way to simplify the problem space. If we de-coupled realism from immersion, and gave up entirely on the idea of reproducing the existing world, we could do more with fewer actuators. This concept would lay the foundations for something new -- the Interaction Centric Approach to Haptic Design.

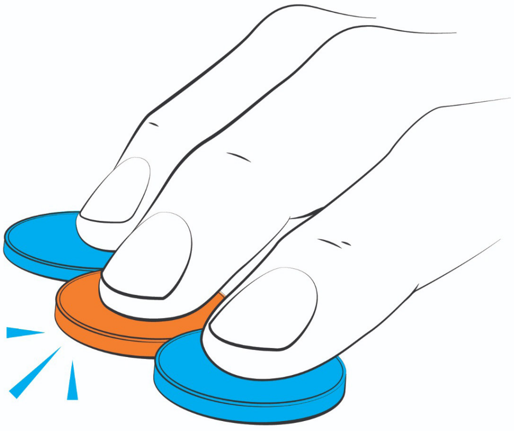

Fig: Hayward's haptic illusion with temperature and three coins.

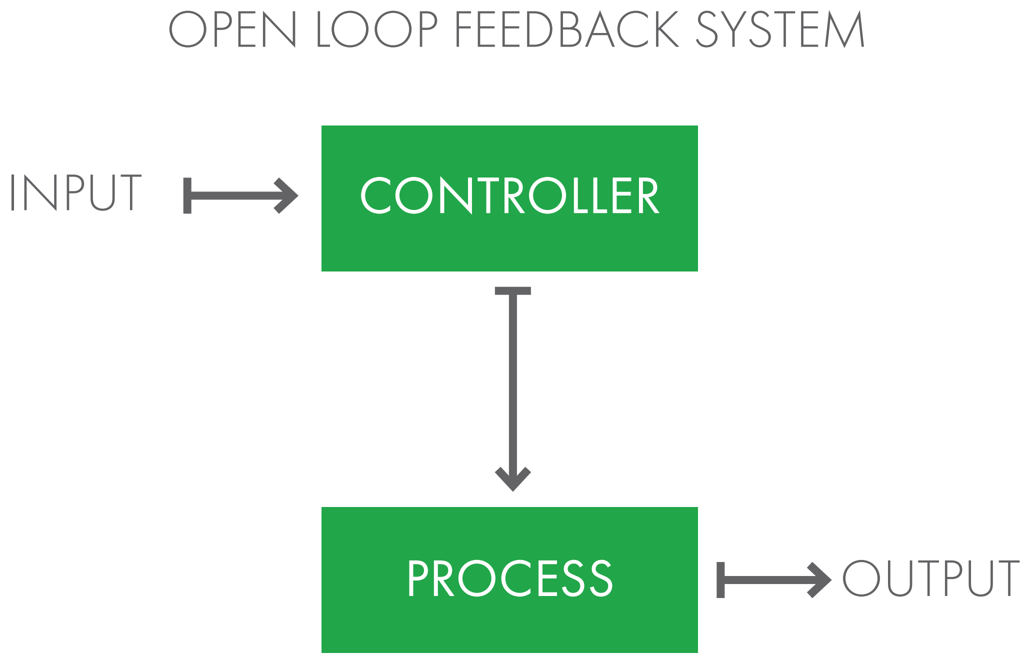

Fig: Open loop feedback system block diagram

THE COMPLEXITY OF TOUCH: A BI-DIRECTIONAL SENSE

Touch is the only sense that allows us to do further interaction and exploring. We don't simply perceive with touch, we can interact and manipulate objects. In so doing we also change the mechanical state of our own bodies/sensors.

This complicates haptic display, as most haptic interfaces have a mix of input and output that blends together. In a force-feedback button, the same finger that provides the pressure input is the detection surface for the output pressure of the button.

When we write our names on paper, the same muscles that put pressure on the pen, and through that, the paper, detect the amount of pressure the paper puts back, keeping us from breaking through the page.

To discuss this, control engineers often talk about open-loop and closed-loop systems. Open-loop systems assume that an instruction, movement, or actuation completes, with no verification. A closed-loop systems senses and measures the

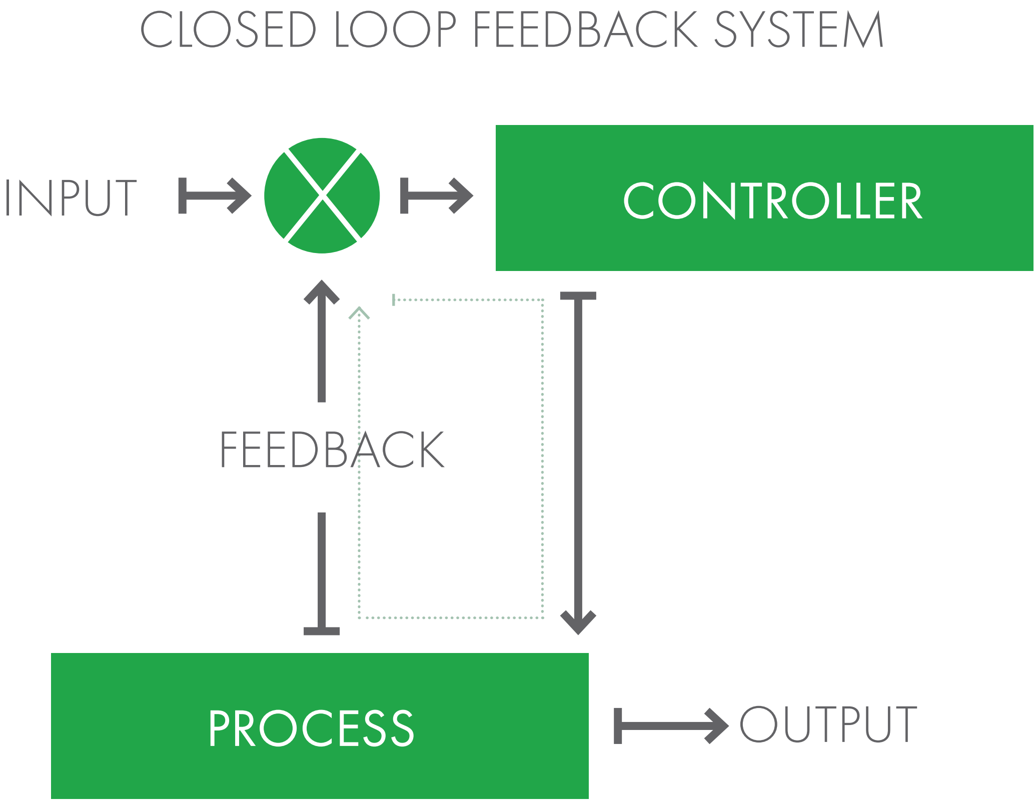

Fig:Closed loop feedback system block diagram. Note feedback loop (dotted)

context and output to modulate the control signal so that the final outcome matches the intent.

Our sense of touch is closed-loop. We're constantly updating our expected perception, and modulating our body to match. When we pick up an object we expect to be solid we hold it loosely as we don't expect the shape to change. If that object instead is a gooey mass, we tighten our grip to ensure it remains within our grasp.

This update is done subconsciously, but the linkage between perceived force and object firmness and even object shape is well defined.[30]

This complicates the nature of designing for touch, as any output/haptic signal has to aim for a moving target. We have to predict not only the perception of the output, but also the user's instinctive compensations to the effect.

CONCLUSION: LIMITATIONS TO THE BODY CENTRIC APPROACH

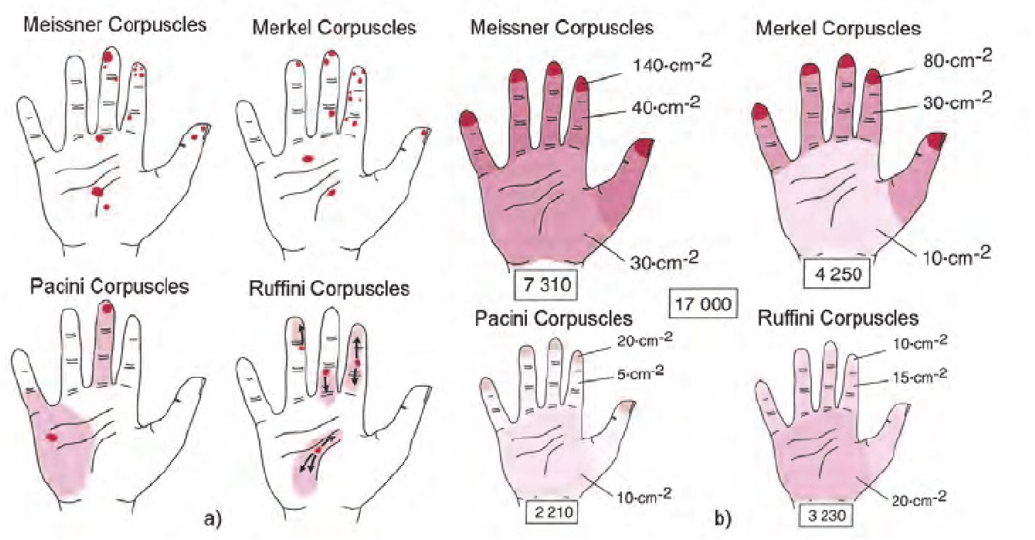

Fig:Density of various mechanoreceptor types in the hands

The body centric approach had been the predominant school of thought in touch for decades, but as a foundation for the work I was attempting to do, it came up short. The body centric approach doesn't build flexible devices, and it doesn't build consumer focused devices, instead it builds towards a research focus. Additionally, from the body centric perspective, I knew what sort of effects I needed to produce, but I didn't know how to produce them. Body centric research rarely detailed the actuator, a key question in my quest to produce an improved controller.

A body centric approach to my challenge -- the Oculus Touch controller redesign -- would lead me down a path that prioritized realism over all other design factors, a tradeoff I wasn't willing to make. Rather than have a controller that worked for dozens of use cases, I would have a product that did one thing extremely well. Furthermore, the body centric approach isolates various sensory channels, and I believed that information integration would be needed. I viewed an IVE as a collection of layers that integrated into a single collective picture. I didn't view the digital world as an exact analog to the real one. Instead I felt that simplified and easily interpretable haptic cues would be more useful -- an idea antithetical to the realism focus of the body centric approach.

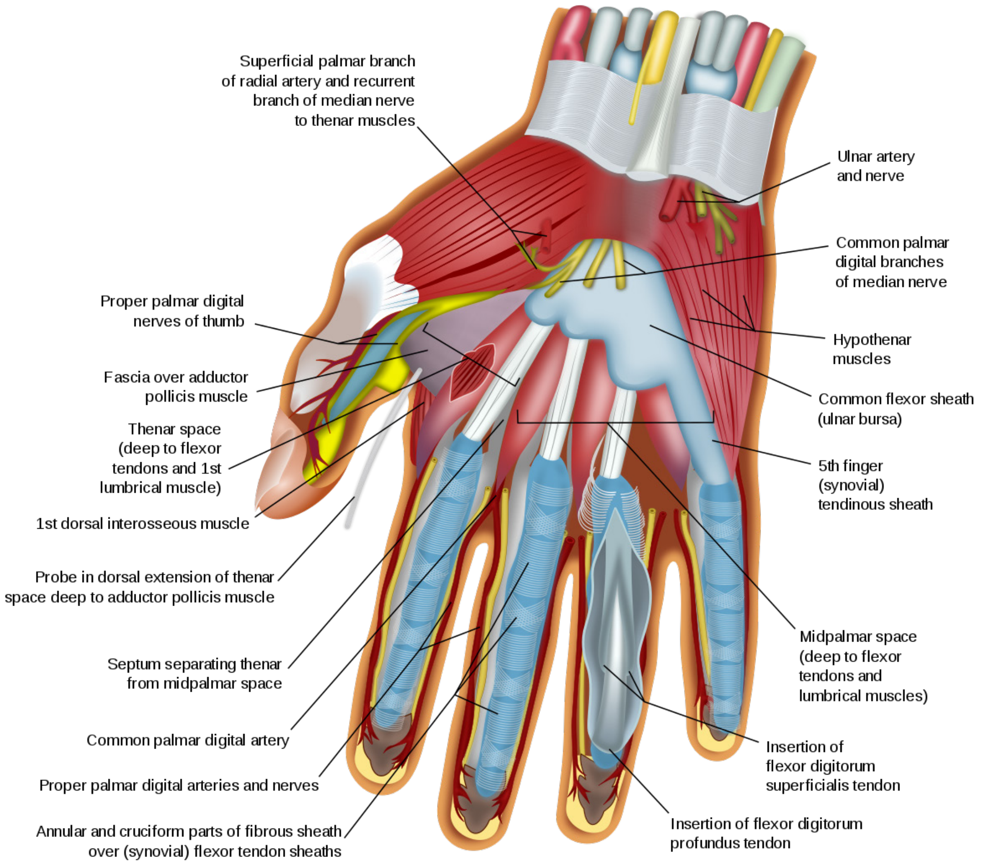

Fig:Location of nerves, muscles, and other sensors in the hand.

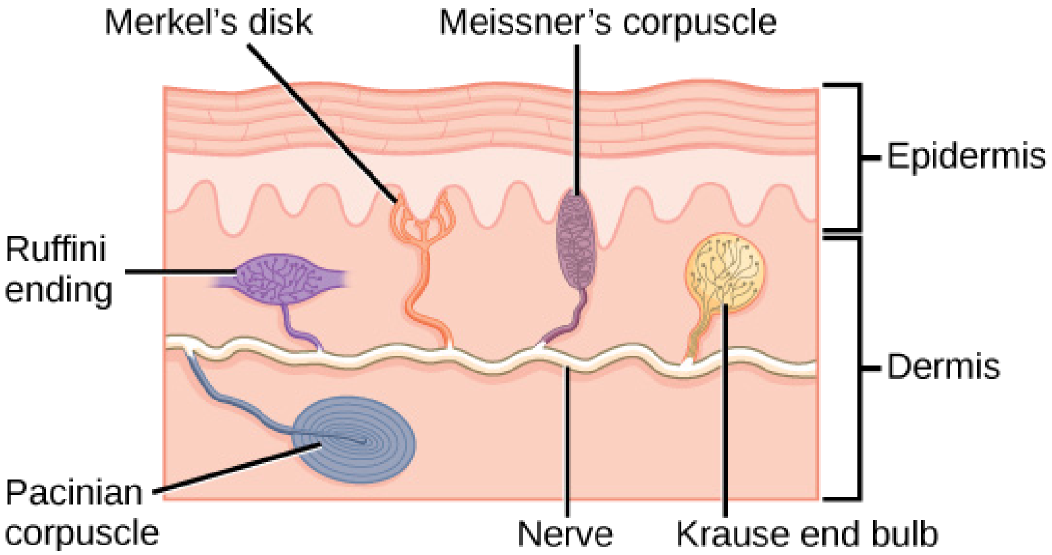

Fig:Cross section of skin and mechanoreceptors.

RESEARCH: THE TECHNOLOGY-CENTRIC APPROACH

CHAPTER 05

THE TECHNOLOGY CENTRIC APPROACH

So if the body-centric approach wouldn't be enough, I needed to explore the other main path -- the technology centric approach. For many, their approach to haptics is driven by the desire to develop and build cool, innovative technology. For these people, there was a focus on the effectors and mechanics that produce haptic interaction, an approach which would become the technology centric approach.

The foundations to the technology centric approach lay in the technology revolution brought on by the integrated circuit/transistor in the late 1970s, early 1980s. While academia had explored the mechanics of haptic interactions long before the microchip, the microchip allowed for commercial enterprises -- Immersion, Aura Systems etc.-- to spring up.

These companies could, for the first time, offer end-to-end solutions. At this point researchers could explore haptics without the need for the teams of engineers required to build new technologies, displays, and interfaces for every test. The ability to quickly develop new haptic motors, controllers and to incorporate them into consumer devices has been one of the most significant catalysts to the field.

The technology-centric explores the "What" of haptics. What sorts of effectors and drivers can be built? What sort of products can be developed with those products? And What effects can be rendered with these interfaces.

CAUSING AN EFFECT: TYPES OF ACTUATION

To recreate the sense of touch, we apply forces, vibrations, or motions to various points on the skin. Haptic feedback exists in our smartphones, video game consoles, and computers, but it also exists outside of consumer goods. The satisfying snick-thunk of a car door is engineered around our sense of touch. However, haptics is tightly married to consumer electronics, where devices are intractable and interactive. If the human/computer interface engagement is one-way, we lose engagement and interest in a product.

> VIBRATIONS ARE FLEXIBLE IN APPLICATION, AND IN INTERPRETATION, AND THUS HAS THE LARGEST DESIGN SPACE OF ALL HAPTIC FEEDBACK TYPES.

These explorations are not new. The Nintendo 64 was the first consumer product to widely implement active haptic feedback, through the now-classic Rumble Pak[31]. The Rumble Pak implemented the most common type of haptic feedback, vibrotactile feedback, but there are dozens of mechanisms that target various tactile-perceptive channels to produce or control

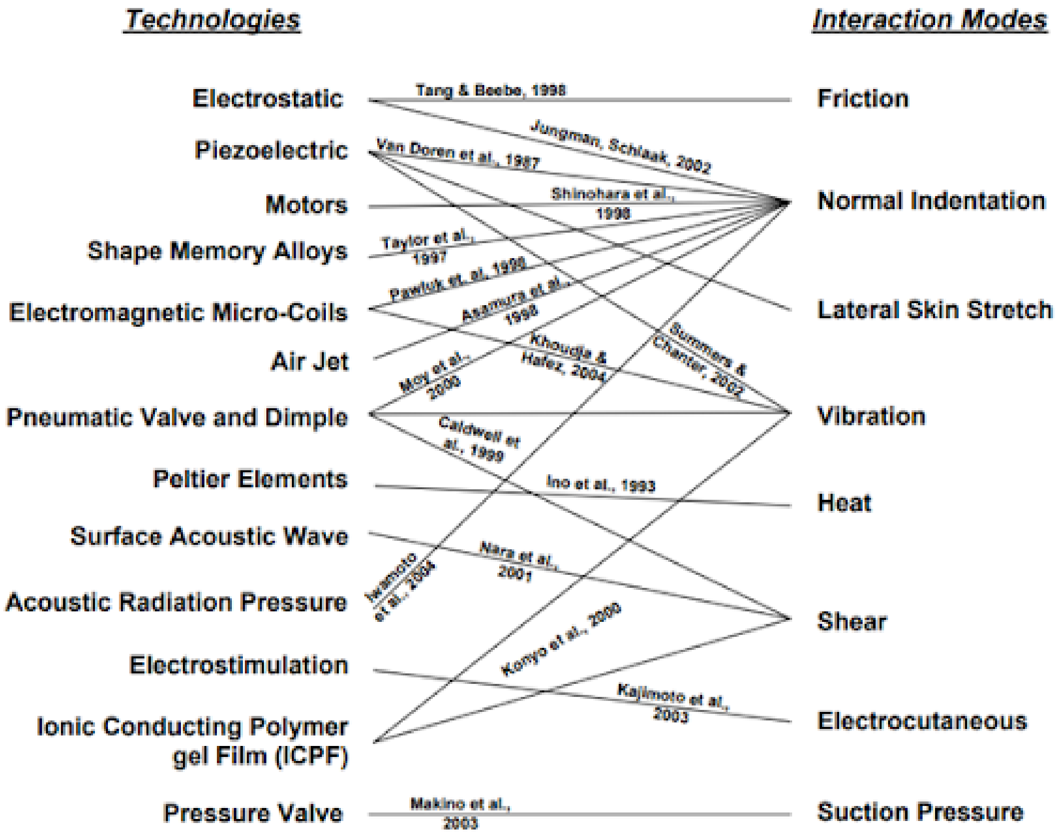

Fig:Map of haptc feedback technologies to interaction modes. From Verrilo (1962)

an effect. The tech-centred approach has also enabled haptic feedback to be used in medical training scenarios via force feedback devices such as the DaVinci Surgical Robot. Each interaction mode or method requires hardware tailored to that interaction. While we can rely on the fallibility of touch (see section "The Complexity of Touch: An Integrative Sense") to help some technologies achieve other interaction modes -- e.g. vibrations can be perceived as heat -- these are the exception not the rule. The figure above outlines over a dozen actuator mechanisms and what sort of sensory system and perception they target.

One of the more useful interaction technologies is vibration. Vibrotactile feedback is easily the most common method of haptic feedback found in consumer devices today. Vibrations are flexible in application, and in interpretation, and thus has the largest design space of all haptic feedback types.

What all the actuation types share, is a need for a signal, something that defines the mechanical behavior of the actuator in terms of output and time. Commonly known as a waveform, this series of position and time points can be plotted to produce a clear chart of motion.

WAVEFORMS, AND HOW TO MAKE THEM

From a mechanical level, there is very little difference between audio and haptics. A haptic waveform is fairly similar to an audio waveform. Many real world interactions produce both auditory and haptic signals across a wide range of frequencies and amplitudes. A car accident generates high frequencies audible that propagate through air, and lower frequency vibrations that are felt through the ground. Putting our name to paper has low frequency contact mechanics and high frequency sliding mechanics. Trying to constrain haptic interactions to a particular bandwidth impacts realism.

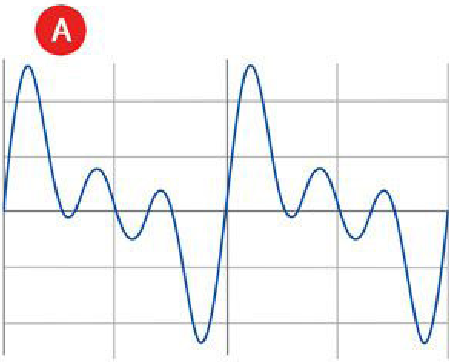

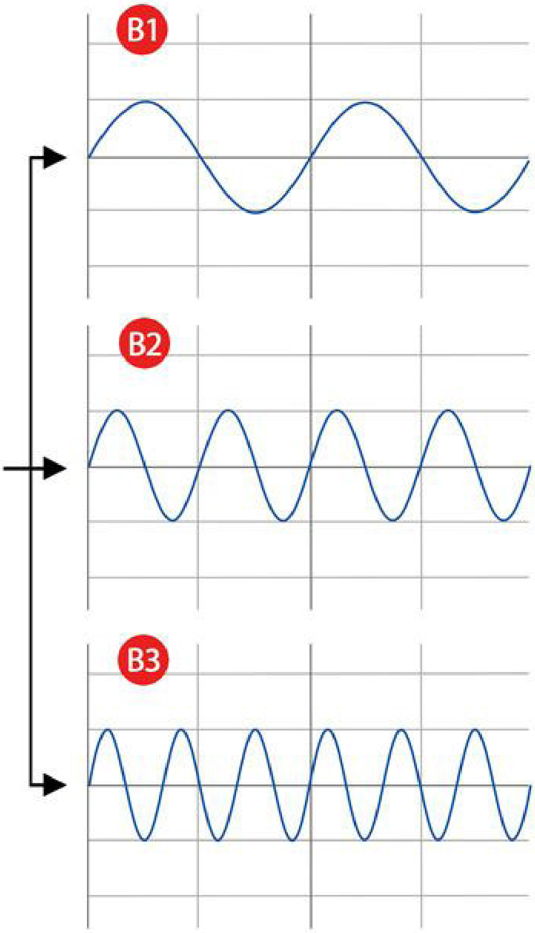

Fig: Complex waveform and decomposition into simple sine waves

So what do these signals look like? Typically, the horizontal (X) axis represents time, and the vertical (Y) axis represents intensity -- measured either by displacement, pressure, or voltage. The figure above illustrates the waveform of the simplest type of sound, a sine wave at three different frequencies. In each sine wave, the oscillation occurs at a fixed rate, or period. By measuring the number of peaks in a unit of time, you can find the frequency. The amplitude, or the intensity of each peak shows the energy contained at that peak. These pure tones rarely occur in nature, auditory and haptic signals typically contain dozens of tones of different frequencies happening simultaneously.

In order to discriminate the various tones intermeshed with one another, we can perform a Fourier transformation,

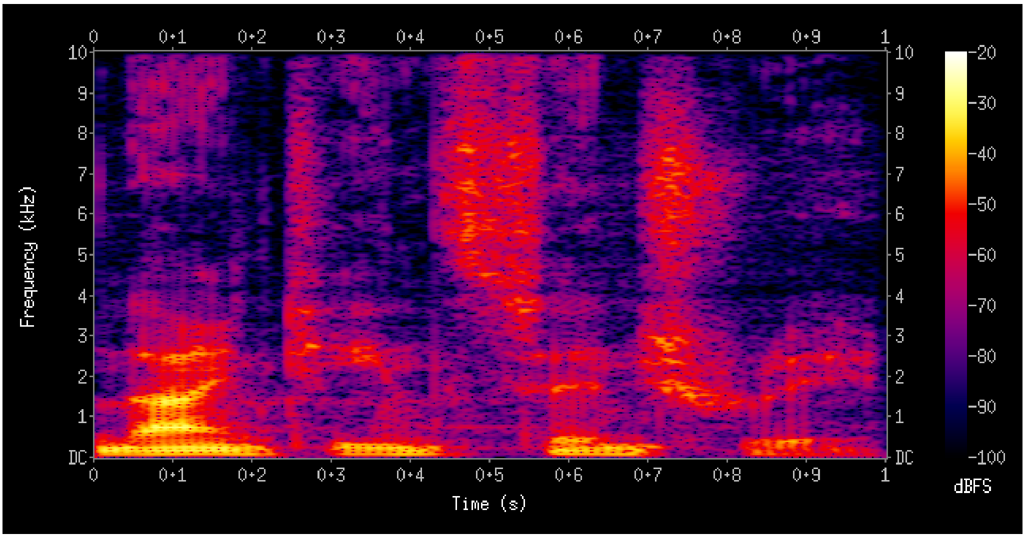

Fig: Spectrogram example

generating a spectrogram -- a chart that maps frequency and intensity. An example of a spectrogram is above, showing time, frequency and intensity of a complex waveform.

To creating a library of reference for each of the interactions we had hoped to recreate, we used a series of recordings of real world interactions. To the right, you can see some of the buttons tested in the "Button Box". These waveforms were created by recording X,Y,Z waveforms through two different accelerometers -- the BNO055 by Bosch, and the ADXL 345 by Analog Devices. Both of these recorders were powered by an Arduino Uno clone. The bottom right shows the ADXL 345 accelerometer mounted to a wooden hand. When taking actual recordings, the accelerometer was mounted to my hand during an interaction

These complex waveforms are constantly occurring in the real world -- when we run our fingers over a surface, our skin deforms and deflects thousands of times a minute.

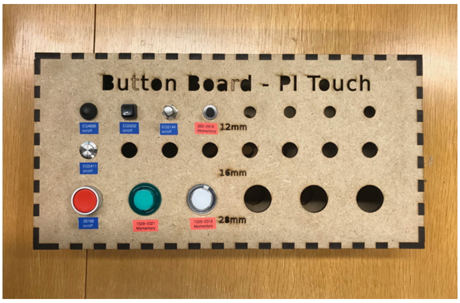

Fig: Button Box

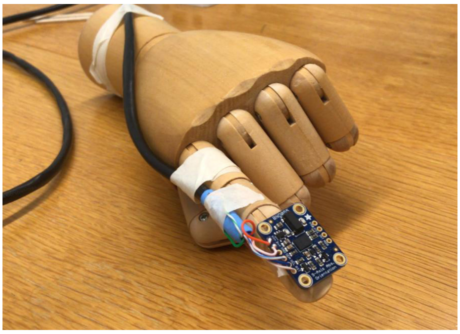

Fig: ADXL345 used to create waveform library.

The difficulty is not in generating a waveform, but in measuring a reference from the real world interaction. There are three main approaches to this measurement.

SOUND-BASED WAVEFORM DESIGN

The first is through the measurement of sound. As mentioned above, the physics of haptic waves are similar to audio waves. The critical difference is in frequency, with the skin able to perceive waves from 0-1kHz (peak 60Hz-400Hz), and the audio range of human hearing running from 20 Hz to 20 kHz, (peak 1kHz and 4kHz). This means that an off-the-shelf microphone can also be used to record an approximated haptic signal.

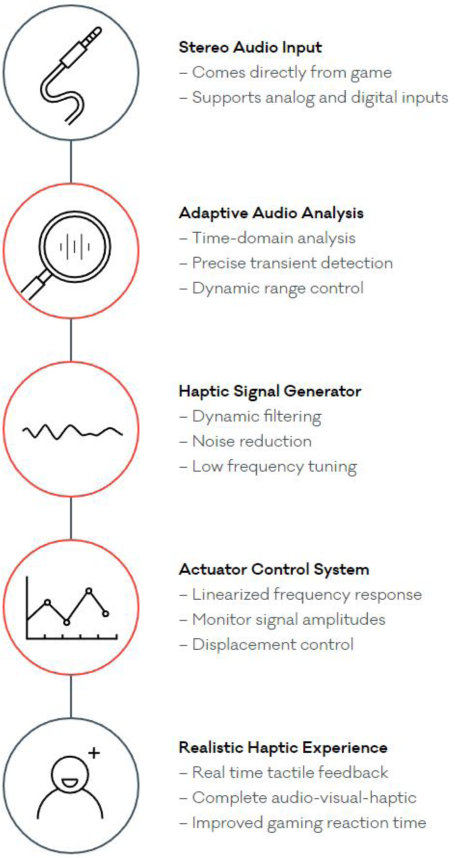

Fig: LoFelt sound based workflow

Place a microphone near a surface and then rub the surface, the sound generated is fundamentally the haptic waveform. You need to filter out the frequencies that fall out of the haptic range. Microphones, however, are not designed for this task, with haptic frequencies sitting well away from peak responses, and near the edge of the useful range. Sound-based haptics is quick and straightforward. Therefore, the sound-based approach offers advantages for IVEs, first and foremost, it can be done with digital objects in real-time. Rather than have to pre-render every interaction, the sound-based approach can automate the process of generating haptic waveforms.

LoFelt, a Berlin-based startup, is developing a Digital Signal Processor (DSP) known as the LoFelt Wave that can pull sound from an IVE and develop an approximation of the haptic interaction from the IVEs audio track. The critical challenge in sound-based approaches to haptic waveform generation is the filtration of the signal. Audio content might include sounds that should not be translated into haptic feedback. IVEs other content can contain additional noise that can cause the actuator to vibrate continuously, diminishing the overall impact of the haptic experience.

In a shooting game, we want to feel the inbound bullets flying by, and the thump of an explosion, as we would in real life. What we do not want to feel is the commands coming in over the radio, or bullets that impact across the map. The audio approach has difficulty in separating these data from one another. Haptic signals and audio signals share parts of their spectrum, and raw audio data does not contain any physics data -- position, acceleration, orientation. This limit makes it hard to extrapolate or modify the waveform, and in an IVE hard to localize in 3D space.

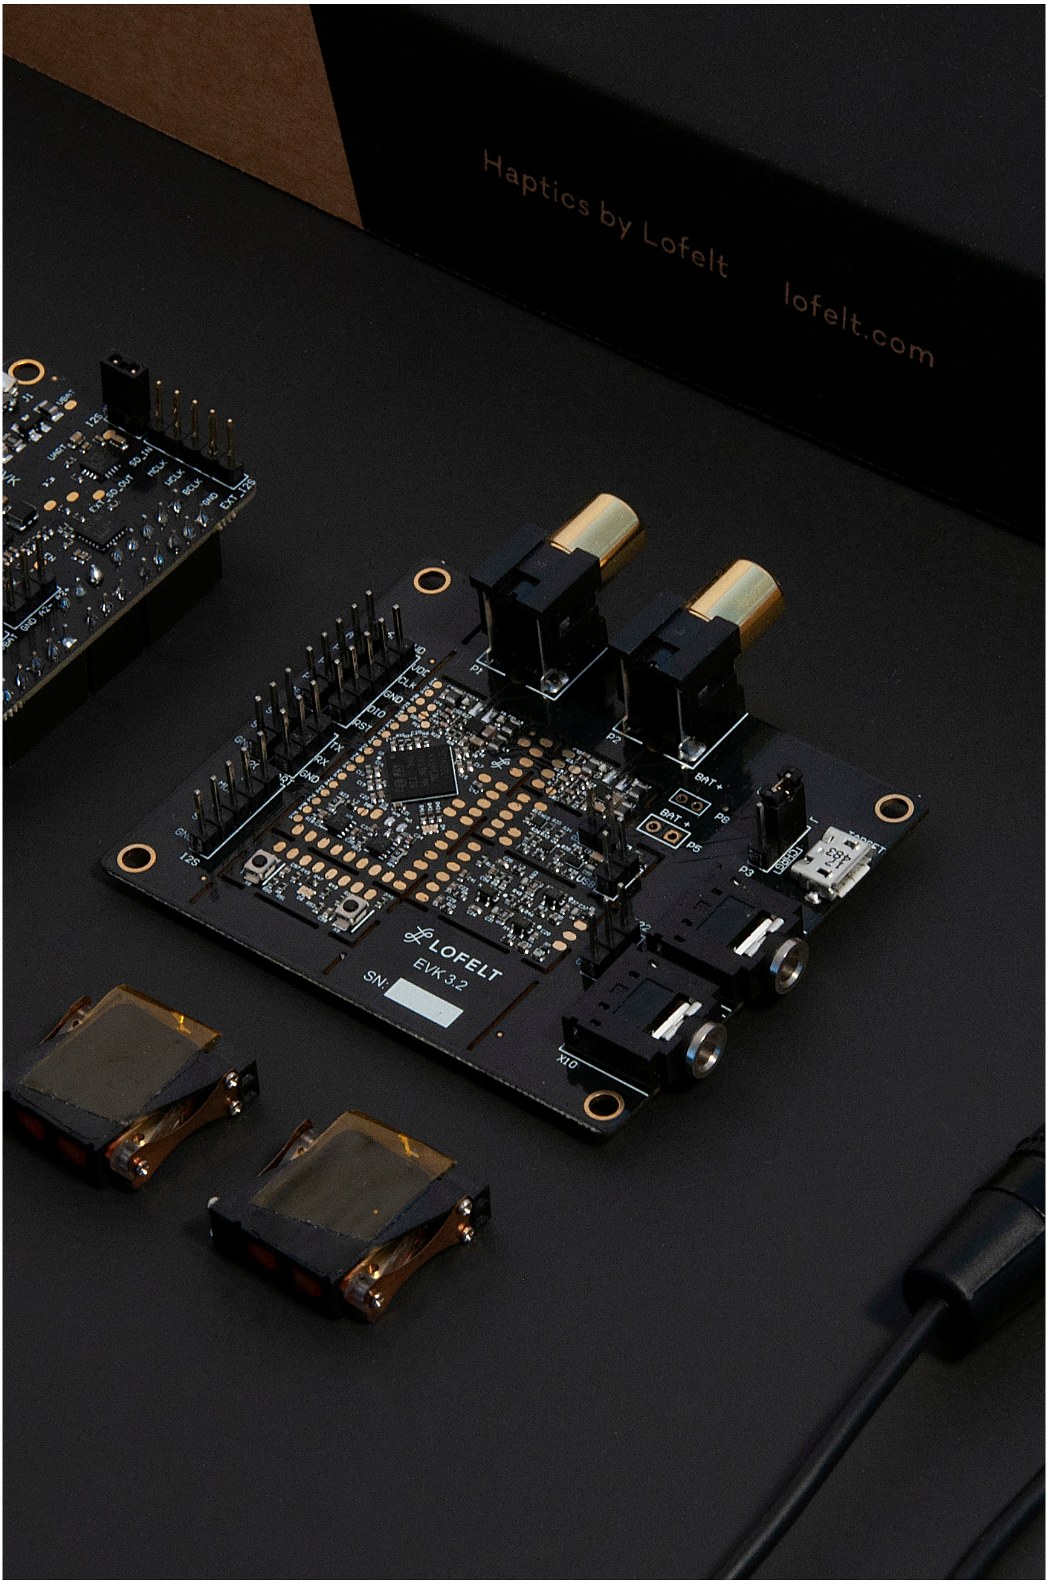

Fig: LoFelt Evaluation Kit (EVK), L5 actuator, running LoFelt wave DSP.

PHYSICS-BASED WAVEFORM DESIGN

We can also reverse engineer a haptic signal from the real world physical behaviors. If a waveform is the mapping of position over time, by monitoring the changes in position, we get a waveform. This approach, the physics-based approach to haptic signal generation, measures deformation (forces) or movement (acceleration) in a probe during an interaction. Ideally, the probe is the finger, or some other body part, allowing physics-based waveforms to come already calibrated for the unique properties of the skin.

If we go back to trying to record the haptic signal of rubbing a surface, the physics-based approach offers two primary tools. The first is through a tribometer, a collection of highly sensitive strain gauges that measure the tangential forces generated by friction. As our fingertip brushes over a surface, the hundreds of microscopic interactions between the peaks and valleys of the surface and our finger generates friction. The smoother the surface, the fewer peaks, and valleys, and the less surface area in contact. The rougher the surface, the more peaks, and valleys, and thus higher surface area. The more surface area, the higher the frictional force. This phenomenon is why skaters apply sandpaper-esque grip tape to their decks in order to stay on the board through tight turns and tricks.

By measuring the microscopic changes in friction as our finger slips over valleys and sticks to peaks, a tribometer can recreate the texture of the surface, and once we know the surface texture, we can generate a waveform that mirrors the peaks and valleys.

We can also explore the same surface interaction by measuring the acceleration of the finger as it skims over the surface. By mounting an accelerometer matched in bandwidth to the human skin on the finger, we can record the microscopic jumps that it makes in X, Y, and Z directions as it travels across the surface. By knowing the changes in velocity (acceleration), we can find the changes in position (velocity), giving us the waveform.

A further subset of the physics based approach is the mechanics a based approach, which also looks at force and

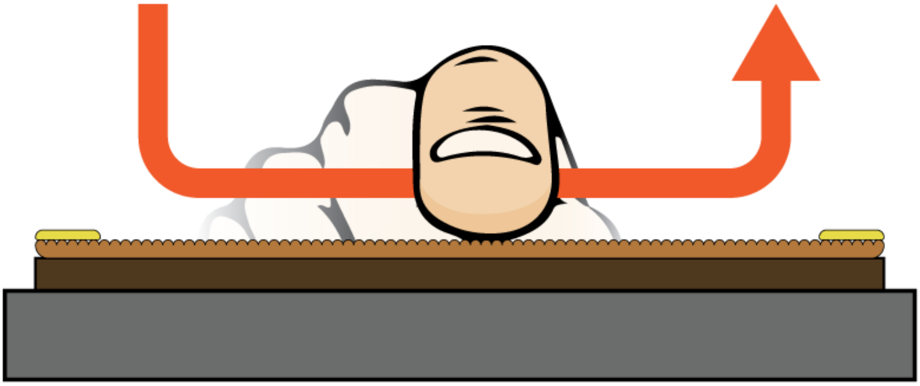

Fig: Graphic depicting how to record an interaction between a surface and a finger.

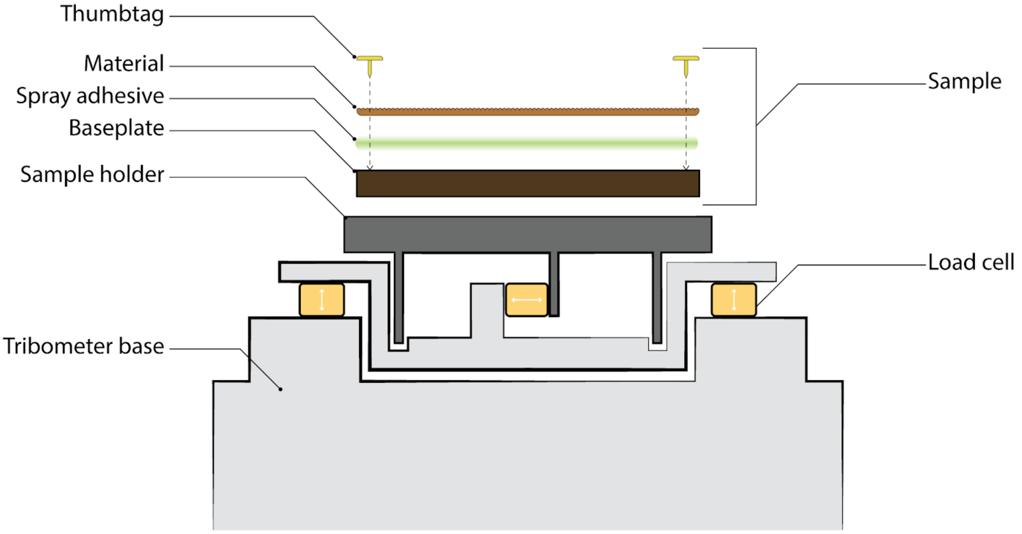

Fig: Tribometer detail.

acceleration data to build a waveform. The critical distinction between the two is that, the mechanics based approach explores the behavior of the object, not the probe. The mechanics based approach takes things a step further than the physics based approach by building a model of the object and generating waveforms from the expected behavior, rather than the recorded behavior.

This approach, pioneered by Kim and Lee [32] was used to build a model of a Cherry MX key switch, commonly found in mechanical keyboards. The mechanics of a keyboard switch are well understood. They first partitioned the force-displacement curve partition the force-displacement curves into a slope section, a jump section, and a bottom-out section. Each section was parameterized, and through the construction of that a waveform could be built.

The physics-based approaches offers the possibility to capture specific interaction conditions and mathematically identify the patterns which are independent of these. For rendering compelling haptic interactions, contact/interaction invariant patterns minimize the need to provide 'online' feedback cues. This is a multi-step process. The physics-based approach can be challenging to pull off, in the digital space, requiring massive amounts of pre-rendered data. What it does offer though is a complete picture of not only the vibrational waveform but also force data -- a critical outcome for non-vibrotactile haptic feedback. What is unique about the physics-based method of recording is that it also calibrates for skin dampening effects, generating a clearer waveform than the alternatives.

AI-BASED WAVEFORM DESIGN

The bleeding edge of waveform generation uses neural networks to generate waveform renderings. The ideal waveform for an IVE is procedural, computing real-time waveforms from IVE parameters. A promising approach is using neural networks to generate data. Neural networks are a set of algorithms, modeled loosely after the human brain, that is designed to recognize patterns. Companies such as Actronika SAS, a Paris based manufacturer of voice coil actuators are actively developing these technologies.

A typical neural network has artificial neurons called nodes arranged in a series of layers, each of which connects to the layers on either side. Input nodes receive various data from the outside world, such as distance, velocity, acceleration, force. On the other side sit output nodes which respond to the information it has learned -- the parameters of the waveform. Between the input nodes and output nodes are one or more layers of hidden nodes, which, together, form the majority of the artificial brain. Most neural networks are fully connected, which means each node is connected to all nodes in the adjoining layers. A number called a weight represents the connections between one unit and another. The higher the weight, the more influence one unit has on another.

However, a neural network requires training data in order to work. The data is used as the benchmark by the AI to determine the success of the nodes' weighting. To generate training data, the AI-Based approach requires data from another method. The success of the neural network is predicated on the quality of the training data, so the AI approach requires a significant initial time investment.

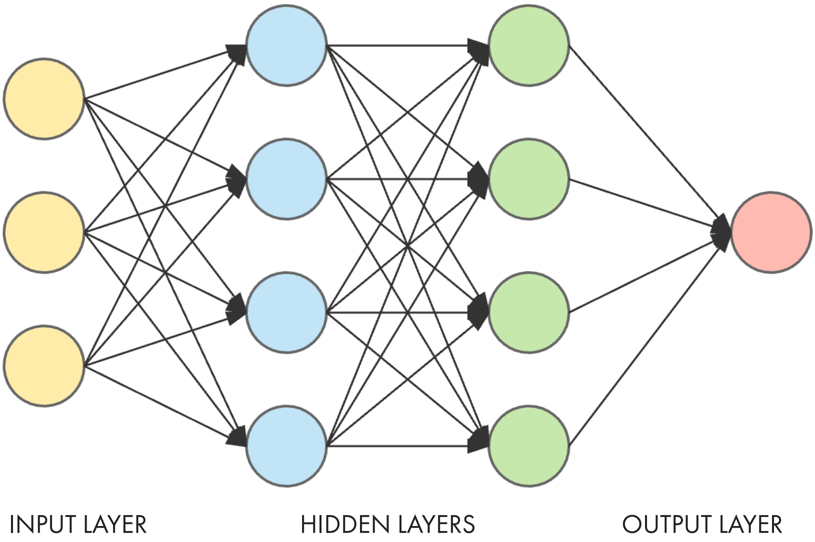

Fig: A simple neural network, with two hidden layers.

The AI approach can be challenging to make work, requiring vast amounts of computational power even to come close to real-time computation. While nowhere near ready yet, an AI approach offers an incredibly flexible approach that may generate extremely accurate haptic renderings. If done in closed-loop, the AI can self-improve the output to help model for variations between actuators and even between physical products.

TYPES OF VIBROTACTILE ACTUATORS

Fig: A render of an Actronika SAS "Haptuator" as used in the inTouch design.

Each type of waveform can be played back on dozens of actuators, each of which has their benefits and their drawbacks. These actuators can be binned into two categories, narrowband and wideband. Narrowband actuators have a small working range, typically a band spanning 10-20 Hz. Wideband actuators have a variable frequency, and thus can target a wide working range, spanning a couple of orders of magnitude. Wideband actuators are much newer and more costly. The following outlines the most common vibrotactile actuators used in products today.

ECCENTRIC ROTATING MASS (ERM

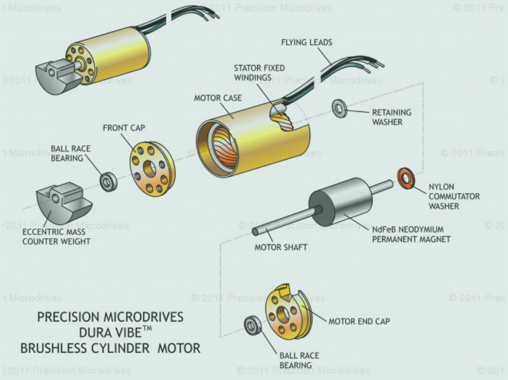

Fig: Exploded view of an ERM

An Eccentric Rotating Mass (ERM) works much like an improperly loaded washing machine -- when mass is unevenly distributed around a point of rotation and oscillation occurs. Larger masses and faster rotation increase the magnitude of the vibration. The ERM motor is cheap and can produce a strong vibration for their size. Most ERMs have a maximum spin of 8000-12500 revolutions per minute (RPM)

However, ERMs have some notable limitations. The amplitude of the vibration sensation is linked to the RPM of the motor. Acceleration amplitude of the vibration grows quadratically with the angular velocity, i.e., the vibratory frequency. Therefore vibrations at a specified amplitude can only occur at a particular amplitude/intensity, and particular intensities of vibration only occur at specific frequencies.

Pros:
\+ Cheap
\+ Easily sourced
\+ Good Power to Weight
Cons
\- Linked vibrational frequency and amplitude
\- Spin up/ spin down times slows response rate

LINEAR RESONANT ACTUATOR (LRA)

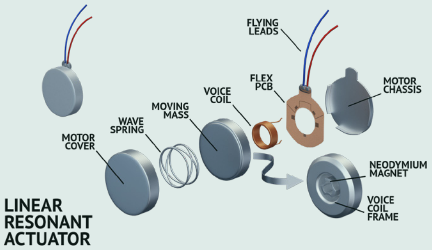

Fig: Exploded view of an LRA

Linear Resonant Actuators (LRA) are the first step up from the lowly ERM. LRAs utilize a moving mass, permanent magnets and springs. A coil generates a magnetic field which moves the mass towards the permanent magnet. The springs dampen the movement and assist with the oscillation.

As with any mass-damper system, there is a resonant frequency to the oscillations, and the closer an LRA is to the resonant frequency, the higher the efficiency. As such, LRAs produce noticeable vibrations in a very narrow frequency band. Thus an LRA has variable amplitude and a fixed frequency. LRAs are driven with AC current alternating at the resonant frequency. This resonant frequency is typically between 200-400hz, coinciding with the peak frequency response rate of the skin.

Pros:
\+ Predictable Behavior
\+ Low Power Consumption
\+ Fast response rate
Cons:
\- Alternating Current power requirements (typically 1.8VAC @ 200-400hz)
\- Narrow operating bandwidth

DUAL MODE ACTUATOR (DMA / LRA+ / 2DLRA)

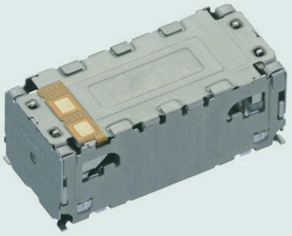

Fig: ALPS Haptic Reactor

Using similar fundamental principles as an LRA, the DMA/LRA+/2DLRA adds in an additional axis of vibration. Newer designs such as the ALPS haptic reactor have multiple oscillation frequencies -- harmonics of one another -- that allow for approximation of complex waveforms by adding in a second dimension of oscillation. By moving in two dimensions at different resonance frequencies (harmonics), the DMA can roughly approximate full waveforms.

Pros:
+ Wide Bandwidth
+ Simple to drive
+ Fast response rate

Cons:
- Unpredictable response range
- Large Size

VOICE COIL ACTUATOR (VCA)

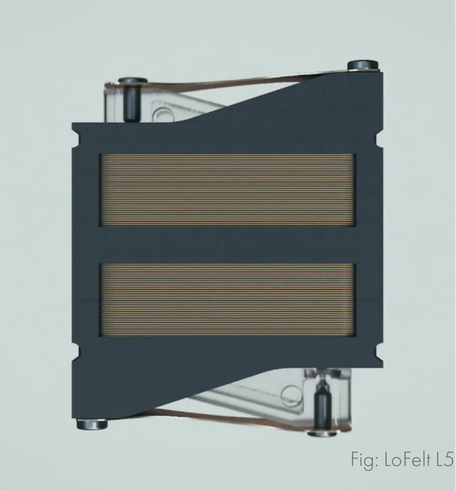

Fig: LoFelt L5

Another improvement on the LRA principle of operation, voice coil designs are just starting to come into their own. By taking the basic principle of an LRA and combining it with the fundamental operating principles of the common household speaker, voice coil designs apply a voltage across the terminals of the motor creating a movement. Reversing the polarity of the applied voltage will move the motor in the opposite direction. The force is proportional to the current that flows through the coil.

The voice coil can run over large bandwidth solving the traditional limitation of an LRA. The voice coil design is difficult to drive, typically requiring filtration of the signal, and has a non-linear response curve across frequency and amplitude ranges.

Pros:
+ Fast Response Rate
+ Large Bandwidth of Actuation
+ Tunable frequency response

Cons
- Complex control schema
- Heat
- Large Size

PIEZOELECTRIC ACTUATORS (PZT)

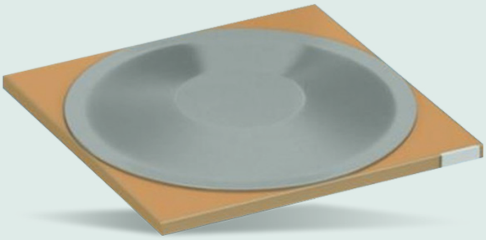

Fig: TDK PiezoHaptic

Piezoelectric actuators rely on the piezoelectric effect -- when a voltage is applied to a piezo material, it contracts. When the voltage is released, the material re-expands to original size. This can be harnessed to produce small displacements and thus vibration. The PZT actuator has an incredibly fast response from a small volume with low mass.

PZT are available in small form factors that enable them to be embedded into tablets and smartphones. However, both require high voltage to drive them, which can present product design challenges. And while piezo actuators and EAPs can work well for finger-tip touch surfaces, they are inadequate for creating the more robust vibrations needed for other devices, such as headphones or game controllers. Uniquely, PZT actuators work as inputs, by pressing on one, a voltage is generated, which can be used as an input signal.

Pros:
+ Can be used as a vibration sensor, not just actuator
+ Large Bandwidth of Actuation
+ Small Size

Cons
- Low Max Vibration Impulse
- Expensive
- Complex Driving Circuitry

MAGNETIC RAM ACTUATOR

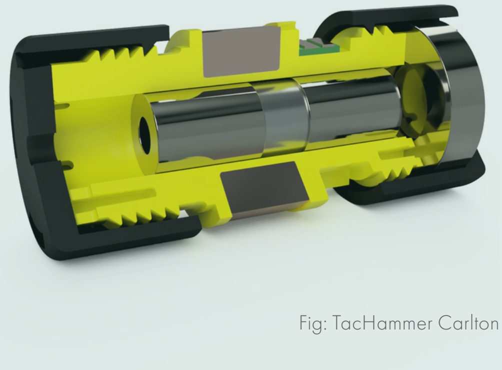

Fig: TacHammer Carlton

Another novel iteration of the LRA, Magnetic Rams combine the core idea of the LRA and a solenoid. To achieve a stronger "click" sensation rather than just vibrate, the Magnetic Ram fires the internal oscillating mass into a material to create sharp strong impacts. The material is changeable and allows for tuning of the sharpness of the impact.

Pros:
+ Simple to drive
+ Strong impact sensations relative to the size
+ Tunable impact sensations

Cons
- Bulky shape
- Weak vibration sensations relative to size.

CONCLUSION: LIMITATIONS TO THE TECH CENTRIC APPROACH

Actuation technology has come a long way from the days of the N64. Video Game controllers buzz, rumble and vibrate in increasingly sophisticated ways. Much of that improvement is due to the technology-centric approach. Indeed, the tech-centric approach does address some of the shortcomings of the body-centric approach, creating a more flexible devices. However, the technology-centric approach can sometimes seem to lack direction. It is the shotgun approach to haptics, designing and building new actuators, and hoping that the actuator can create the right set of experiences for whatever products contain it.

I had hoped that my background and experience in mechanical engineering would align nicely with this approach. The technology centric approach put the actuator technology in the forefront, exploring new methods and mechanisms that have helped revolutionize the way in which we design and develop comprehensive IVEs. While I would draw from this approach, throughout my design process, it wouldn't meet my needs. The technology centric approach led to haptic designs that were poorly contextualized, often mismatching the haptic feedback requirements of the product and the haptic feedback possible within the product. This mismatch led to noisy, immersive signals that overwhelmed or underwhelmed.

Fig: N64 RumblePak

Product development that derives from the technology centric approach is often akin to slapping some vibration motors into a product -- it's viewed as an item on a checklist of "components this product requires". While this approach may work for cell phones or current generation consumer goods where the goal of haptics is to notify the user rather than to recreate missing sensory cures; it falls short in the development of products that interact and support IVEs. Comprehensive and compelling IVEs require more. These IVEs require a new approach to haptics that provides direction and structure to actuator implementation rather than simply actuator design.

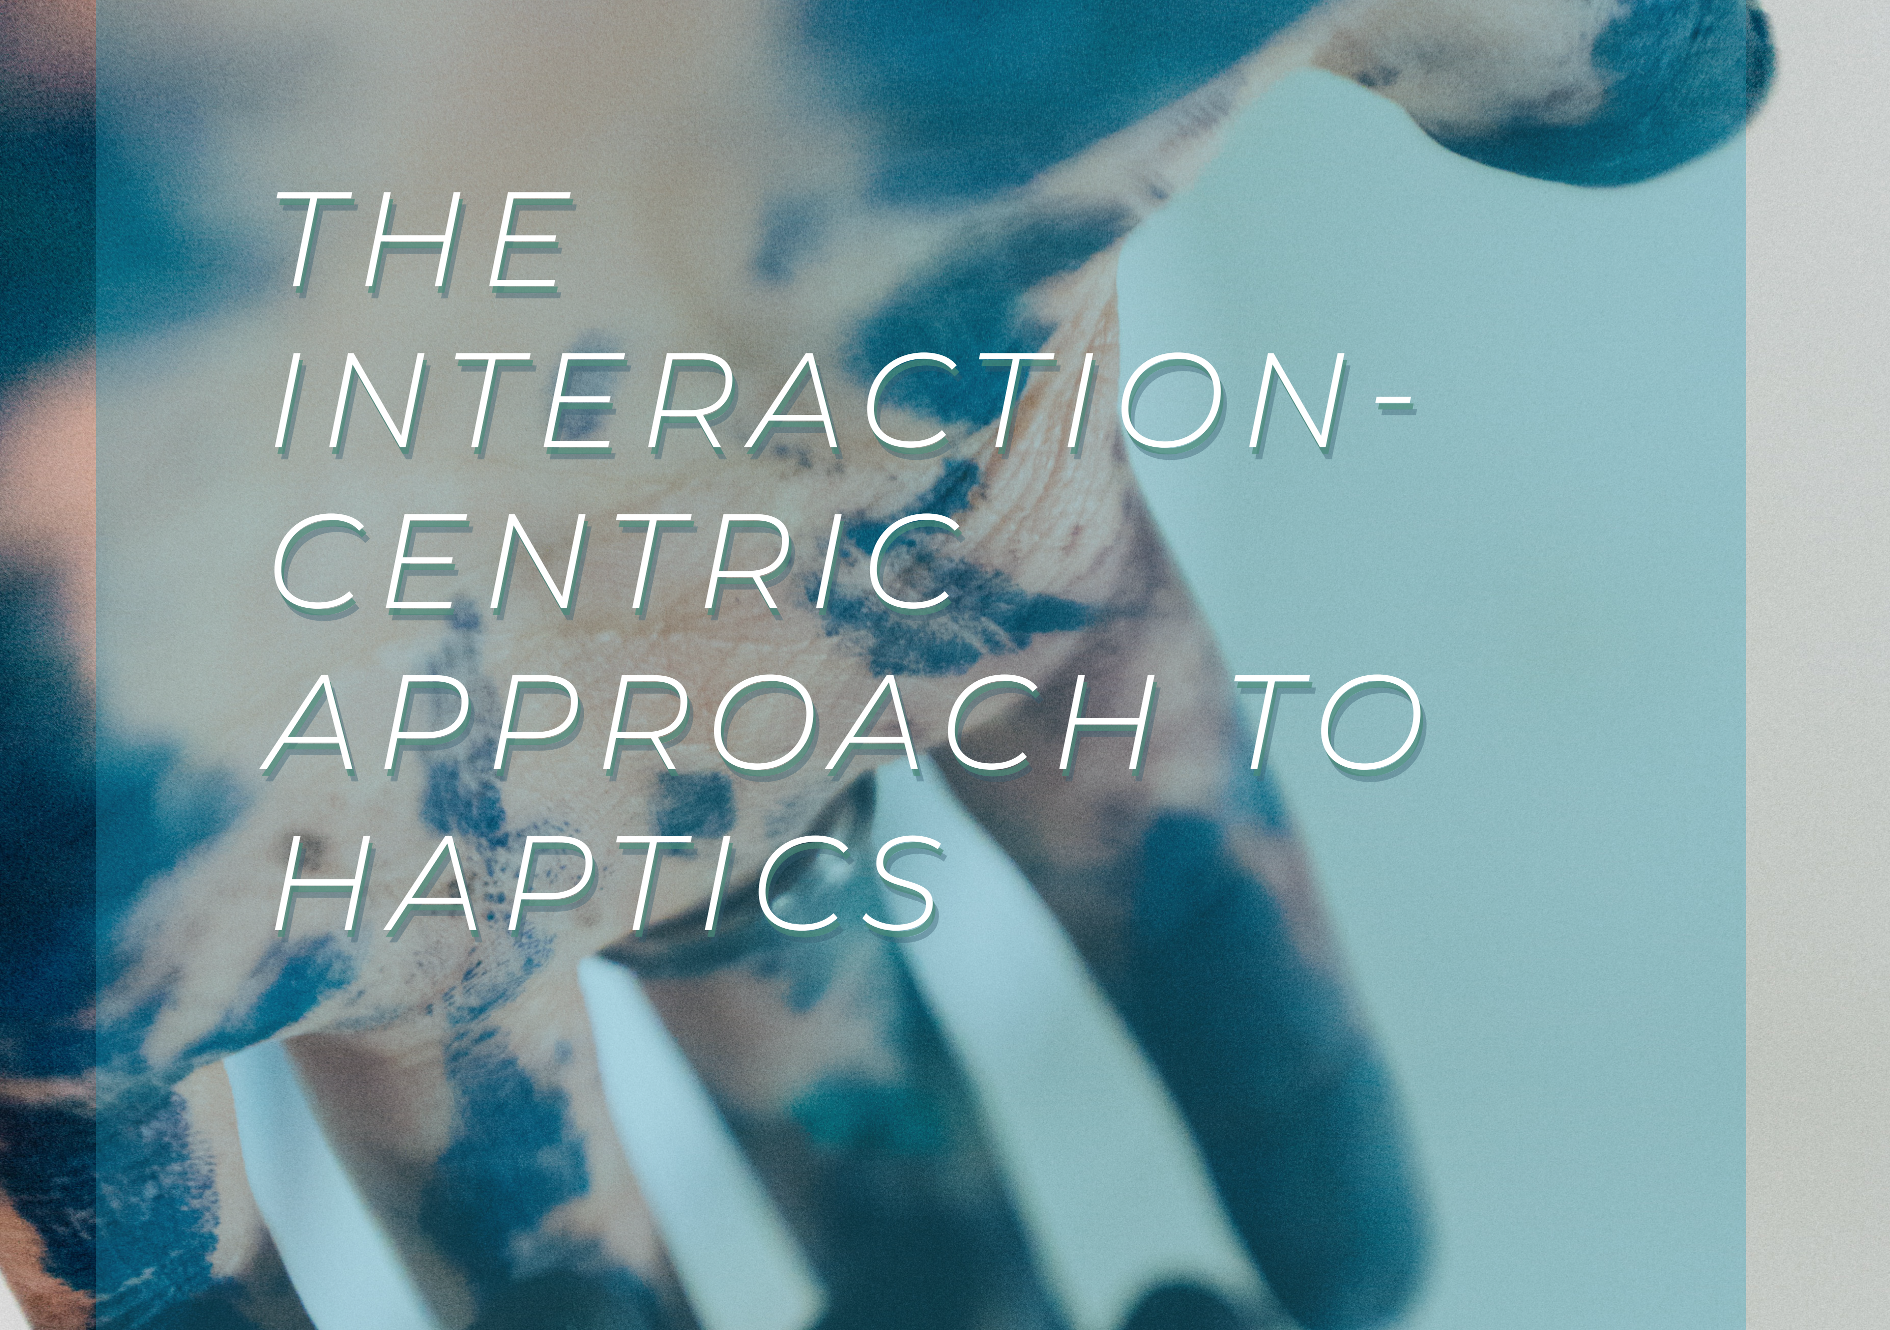

THE INTERACTION-CENTRIC APPROACH TO HAPTICS

CHAPTER 06

THE INTERACTION CENTRIC APPROACH TO HAPTICS

Long before I began this project with Oculus, I felt there was a gap in the approaches used by academics and engineers alike. Surely there was a way to explore haptics from the perspective of the experience. Rather than choose between something targeted and clinical (body-centric) or scattershot and unpredictable (technology-centric), I wanted an approach that was flexible but directed. I wanted to explore haptic interactions not from the lens of psychophysics or engineering, but through information theory -- understanding what was the critical information that needed to be transmitted to make the interaction feel compelling. Rather than answer the "What" or "How" of haptics, I wanted something that solved the "Why" -- why is the haptic channel needed to create a compelling experience? In short -- why bother?

To me, the logical place to begin was the interaction. If the goal was to create compelling and immersive experiences, then the interaction should be the forefront. The interaction centric-approach had to answer three main questions:

› What is the information that makes an experience compelling?
› How much of that information is haptic vs. contextual?
› Why is the inclusion of haptic information critical to an experience?

HAPTIC INFORMATION THEORY -- HOW TO DECONSTRUCT AN INTERACTION

In order to answer the first question, I needed to explore the information presented in an interaction. At its core, interaction is a transfer of information from one being or object to another. The act, and the response or reaction, together comprise an interaction. The critical question in understanding an interaction is what information is passed along? To answer this question, we have to pull apart the interaction, breaking it down into smaller component bits. To help understand this, I'm going to break down an ordinary and everyday interaction -- manipulating a pen when signing our names.

THE EXPERIENCE

First, we look at the experience. This element, the largest, includes an interaction; and the spacial, emotional, and temporal context of the interaction. The purpose of interaction-centric design is to produce compelling experiences, and therefore designing a product from the interaction centric perspective allows us to take advantage of this contextual or 'excess' information, tailoring our approach around the context to help set the stage for the desired experiential objectives.

When we sign our names, there are typically substantial ramifications and risks involved. We sign our names on bank paperwork, health insurance information, and contracts -- all legally binding documents to show our promise to adhere to the terms in the text. Our signatures are tangible imprints of intent, and therefore signing our names on a document can be a nerve-wracking experience. Thus a crucial question to explore at the experience level is what sort of document are we signing.

The answer to this question helps us understand the emotional context behind the interaction. Signing our names on a greeting card carries a very different intonation from leaving our signature on loan paperwork. By defining this emotional circumstance, we now have a design direction for the experience we want to create, the overall objective for the entire exercise.

THE INTERACTION

With 'design information' drawn from the experience, we can now move on to the interaction. At the interaction level, we explore the bulk movements involved. How does each participating user or object -- "player" -- act during the interaction. What are the conditional dependencies between the players: How does each player react to the preceding actions. We can also begin to explore the anticipatory behaviors -- how each player prepares itself for the next step in the sequence. Here we create the storyboard -- the sequencing of action.

In the interaction of signing our name, there are three participants -- the writer, pen and paper. The writer is the only active player, with the other two as passive. The interaction starts when we pick up the pen and finishes when we put the pen back down. The writer anticipates the contact with the paper, slowing down their movements before contact. The storyboard of a signature is relatively simple, with a handful of discrete events, outlined below.

THE EVENTS

If the interaction describes the entire sequence of movements and behaviors, the next levels down are the discrete action/reaction pairs. In this level, we began to look at the physics, with each action, or event, having a corresponding cause. The events are the active sub-component of the interaction, the causes to the effects.

In writing our signature, there are several discrete events. First, the writer picks up the pen. Then the writer puts pen to paper. The pen then transfers ink to paper while the writer moves. Finally, the writer withdraws the pen and sets it back down.

THE EFFECTS

The effects are the reactions to the events. Each event causes a biological or neurological response. In the effects layer, we begin to pull from the body-centric approach to haptics. We have defined the inputs

THE EFFECTS (CONT'D)

to our body -- the events -- and now need to look at the outputs -- the effects.

The first event is picking up the pen. When we pick up a pen, the effect we feel is the weight of the pen in our hand. Another effect is the force we put in hold the pen in place.

The second event is when we put pen to paper. The first effect in this event is feeling the pressure of the pen tip to paper -- in a clicky style pen, we also experience the spring's loading as it is compressed. We feel the paper deform from the spring.

The next event is as we move across the paper. In this event, we feel again the effect of the paper deform. Additionally, we experience the effect of the texture surface texture against pen. We feel the ink flow, changing the frictional forces and lubricating travel. We feel our muscles tense and relax to produce the smooth curves of our name.

As the pen withdraws, the load transfers from paper to hand, and we feel the weight of the pen again in our hand.

THE PERCEPTIONS

In the perceptions layer, we transition to information theory. When exploring the perception, we hope to answer what is the critical information sent between participants. Here we navigate the balance between two design factors -- realism and immersion. When designing a product, it is easy to convey haptic information, but often the haptic "signal" is diluted by sensory "noise," miscuing the user's brain.

Realism adds signal, increasing the risk of noise. A realistic product is akin to watching an orchestra, if totally in tune, on key, and without distractions from the audience, it's easier for the listener to get completely engrossed in the performance. However, if anyone is offbeat, or the audience is talking, the experience is lost -- we focus not on the music, but the distractions -- the noise. An immersive product works the other way; we sacrifice perfection for an increased likelihood of achieving good enough. Rather than an orchestral score, the interaction is an iPod and some earbuds -- the sound is less rich but more flexible and repeatable.

In perceptions, we have to make the same tradeoffs. For the act of picking up a pencil, how accurate does the weight have to be? Do we even need to create the perception of weight, or does an abstraction such as hand-presence work? Does the finger position matter, and do we need to recreate the hand orientation? When we write, do we need to feel pressure, or is the vibration over paper enough information? In this level, we're faced with the first time constraints of the product, and thus the requirements.

THE WAVEFORMS

Once we have defined the goal of the experience and the critical information to recreate that we can start to select actuators. With the perceptions well understood, we can begin to determine what sort of technologies can be used

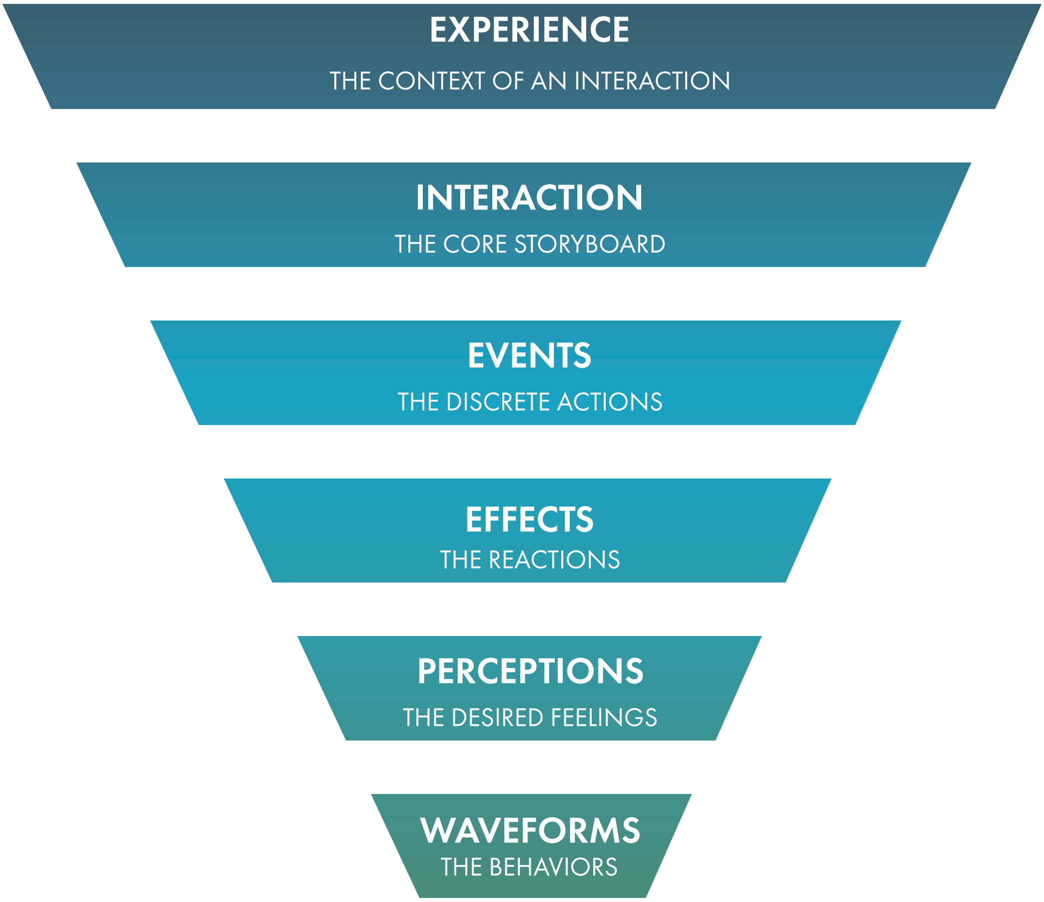

Fig: Haptic Deconstruction

to create an experience. The decision to use a particular actuation, and how that actuation is played back come together to create the most foundational element of the experience.

For each perception we want to create, there is a corresponding actuator. For the signing of our name, we require an actuator or set of actuators that can replicate the bump of contact with the paper, the vibrations of the pen stick-slip, an actuator for frictional force (if needed), and an actuator to produce weight.

For waveform, we need to know the amplitude, frequency, and time. For a haptic cuing approach, we have more flexibility in amplitude and frequency than if we use a more realism focused approach. Realism will require skin stretch and kinesthetic actuators, while a haptic cuing method will need stronger temporal control to link the more abstract effects of vibrations to visual and auditory cues.

We can recreate waveforms for this interaction using just an LRA and relying heavily on visual cues, or use skin stretch actuators and exoskeletons in conjunction with voice coil actuators to produce more realistic and accurate haptic effects.

DEFINE THE CONTEXT

To answer the second question behind the interaction centric approach, I would have to look at the various sensory channels available to the designer. My primary assumption in developing the Interaction-Centric approach was that touch rarely occurs in a vacuum. Reality is a multi-sensory experience. As we navigate the world, we interact with all our senses -- experiencing the sights, sounds, tastes, smells, and textures of everything we encounter. It is this sensory richness that makes life vibrant and defined. Indeed, when we remove individual senses from this experience using anything from a blindfold to a sensory deprivation chamber, we feel uncomfortable, off-balance. Without the full depth of sensory information, humans are on edge.

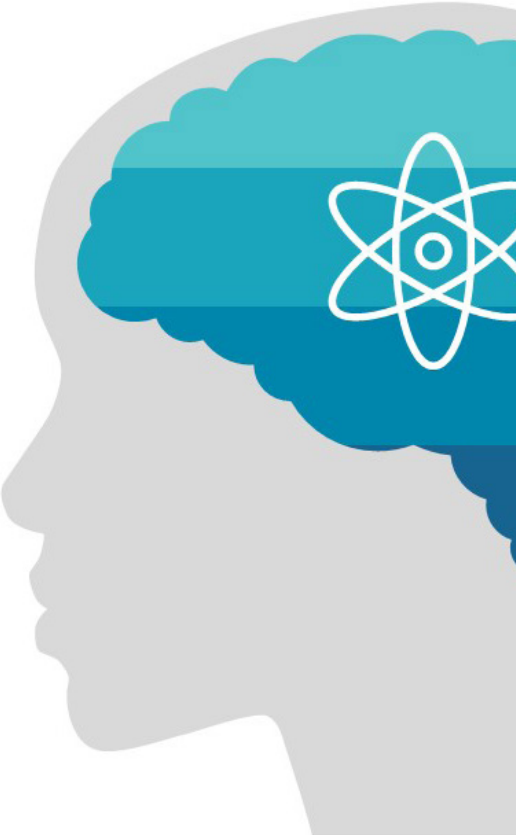

Immersive virtual and augmented environments aim to recreate this lush, rich sensory experience. Indeed, the promise made by companies such as Oculus is to create a virtual reality -- a sensory experience akin to true reality. Their products aim to create alternate worlds we can explore, navigate, and travel.

An IVE is a simulacrum of truth -- an illusion or magic trick on a user -- tricking them into believing a new reality. To do this, an IVE must stimulate our senses in just the right way, showing the user what they need to see and hiding away the sensory noises that conflict with this illusion. This idea -- haptic cuing -- explored the IVE as a balance between signal and noise. Haptic interfaces can convey haptic information, but often the haptic signal is diluted by sensory noise, miscuing the user's brain. Our brain is wired to recognize patterns which fit with our expectations/predictions about the world -- we are hypersensitive to things that feel wrong, the small issues that set our hair on end. This means that any stimulus presented in an IVE must be carefully thought out and contextualized. Anything not totally necessary to building a more comprehensive experience must be stripped out, and all stimuli supported by other sensory data.

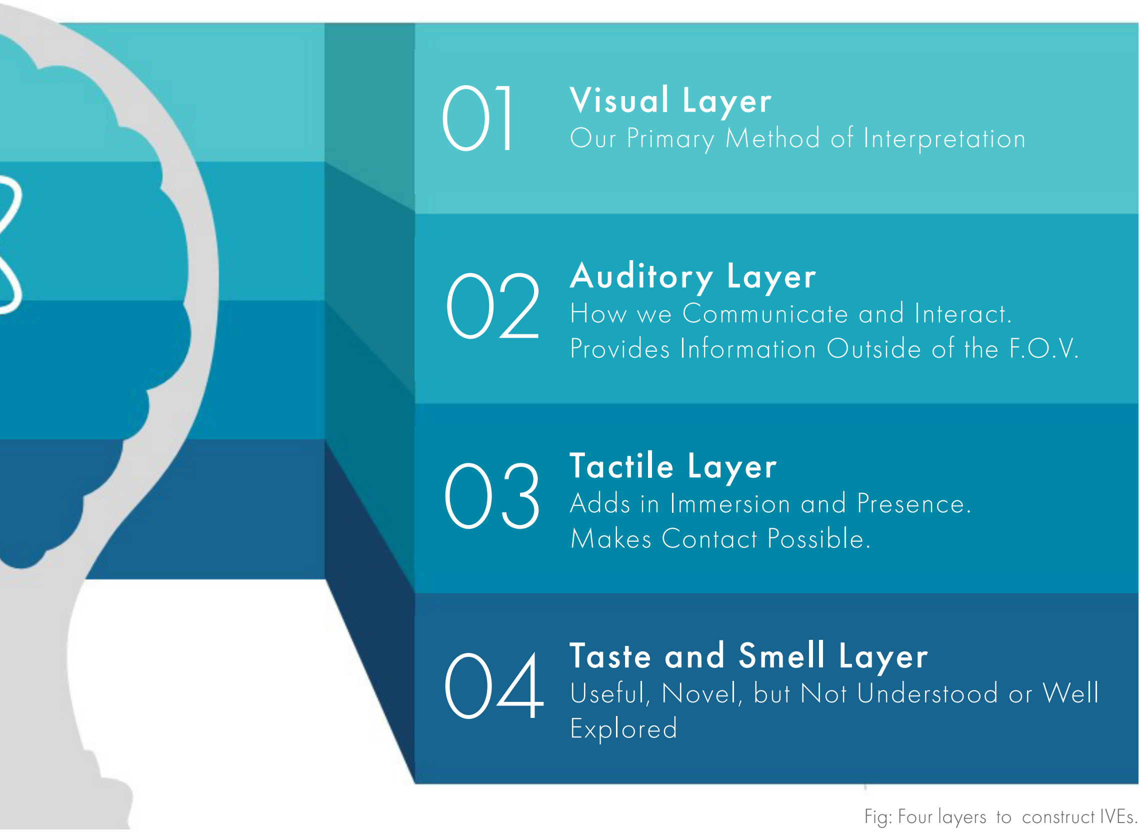

Fig: Four layers to construct IVEs.

In doing so, we can build an illusion that is multilayered, with each layer interacting with and supporting other senses, thus increasing immersion. Without the utilization of multiple sensory layers, these worlds feel flat. Explosions that "THUMP" feel more compelling than a simple flash of light. Car races where it is possible to feel the tire clip correctly clip the apex are more rewarding than a visual pop up. Haptics is best in context, with each sensory strand weaving together to produce a compelling, immersive tapestry.

LAYER 01: THE VISUAL LAYER

If a compelling IVE is a multilayered illusion, the first layer is the visual layer. As adults, we rely primarily on our vision to navigate and interact with the world around us. The visual layer is the foundation upon which we build IVEs.

Presented through screens, projections, or increasingly HMDs such as the Oculus Rift, the visual layer does the majority of the legwork in presenting the IVE. On the visual layer, we find

THE VISUAL LAYER (CONT'D)

the colors, shapes, that the brain processes into objects.

The visual layer has two sub-layers, the foveal, and the peripheral. The foveal visual layer covers the core of our vision -- the central high definition point of our field of view. This cone, +/- 20 degrees off-boresight, encompasses the sections of our vision that are in focus. The peripheral runs from the foveal to +/- 80 degrees off boresight. Each layer plays a different role in perception, with fovea data being "high definition" and the peripheral being "low definition" but more sensitive to movement.

Our eyes take in light rays through the cornea, the clear "window" in the front of the eye, and focuses the rays to the back of the eye, where two types of sensors, the rods, and cones, sit. Rods are responsible for scotopic vision -- the perception of dark/light contrasts, and work in black/white. The rods have a low spatial acuity and sit in the peripheral sub-layer. The cones are what identify color data, and exist primarily in the fovea. Cones detect waves between 430-770 THz, a wavelength of 380-740 nM [17], or more colloquially known as the rainbow. The fovea is the information-dense part of the visual layer, and new HMD screens try to target the fovea with higher definition images. To increase rendering speed, new HMDs track the eye and present higher quality information to the foveal field of view.

The visual layer is built then from color, shapes, motion, and contrast information. Each of these combines to create the world we see, with memories contextualizing combinations of shape and color to produce objects. The higher the quality of the visual display in an IVE, the more immersive the environment. An immersive and compelling IVE relies on the visual layer to drive the experience, and carefully balances data between fovea and periphery to lay the foundation upon which additional layers can stand

LAYER 02: THE AUDITORY LAYER

The second layer, the auditory layer adds in sound. While the visual layer may be the primary driver of content, without the right audio cues to match the visuals, the brain doesn't buy into the illusion. The auditory layer can add additional depth to a virtual world -- e.g., augmenting the imagery of virtual grass with the gentle "swish" noises, or giving explosions the much needed "BOOM." There is a reason that comic books rely on the onomatopoeic "kapow" and "wham" when illustrating fight sequences -- our imagination requires the sound effects. Without a convincing set of sound cues can demonstrate actions or events, visual layer stimuli will never be immersive or convincing.

The auditory layer also slews our sensory system onto a target. The sounds of the world alert us where and when to look at things. We are aware of movements hidden behind walls and barriers at a distance due to the sounds they make. If an action occurs outside our field of view, the auditory layer presents us with the majority of the available data.

The auditory layer also helps isolate us into an IVE. Sound canceling technologies, or even merely drowning out background noise helps remove conflicting elements that do not match the IVE. Acoustic information can ease users into an IVE, either through the presence or absence of sound. Indeed, would the “Star Wars” scrolling text introduction be as effective in introducing us to a galaxy far away without the signature opening chords?

Sound moves us. Music has long been linked to emotional states, and various frequencies have been shown to stimulate brain state. Sounds cue us to the emotional state expectations of an IVE. If we go back to Star Wars, would Darth Vader’s first entrance be the same without the Imperial March? Would we have the same combination of fear and awe? Love songs are much the same -- with Sinatra’s crooning “setting the mood” for romance. In an IVE ambient sound (or the absence) can drive fear in a video game.

Dead Space, a space-based horror game, is known for their sound design. In an excellent example of the audio layer’s power, Isaac -- the player controller protagonist -- exits the confines of the ship and enters Zero-G Space for the first time. This is one of the quietest moments in the game; players can hear Isaac’s breath and the clank of this boot heels against the hull of the ship, but little else. The very first room Isaac steps into when he reenters the ship is a large mechanical bay filled with flashing lights and an overwhelming cacophony of metal clanging. This change from eerie stillness to overwhelming chaos gets the heart racing, with every nerve on edge expecting an alien attack.

The auditory layer is also what enables communication. It is through the auditory layer that IVEs become multi-user and collaborative. While the auditory layer can set the state for interactions, communication happens through words and sounds. Communication is for most of us, synonymous with talking. In an IVE, we issue commands to our comrades, plan strategies, and collaborate through verbal communication. Without the audio layer, the IVE isolates instead of bringing users together.

LAYER 03: THE TACTILE LAYER

The third layer, the tactile layer is the main driver for immersion and presence. The tactile layer is what makes virtual worlds interactive, allowing us to reach out and grasp, and manipulate objects in the IVE. Thus, the ability to convey accurate haptic information -- particularly (vibro)tactile -- is a critical barrier in creating comprehensive IVEs. The primary advantage of haptic feedback is that it allows users to touch and feel virtual objects, helping to eliminate the experience of a ‘false’ reality[18].

To engage our sense of touch, haptic displays convey haptic details, -- shape, size, surface roughness. In each tactile interaction, these features help users identify an object or material. When perfect, such determinations are possible; unfortunately, the haptic “signals” presented through

THE TACTILE LAYER (CONT'D)

such displays often differ from the real-life equivalent, miscuing the user's brain. These differences produce "noise" in the signal, diluting the perceived realism and immersion. The higher the level of distraction, the more muddled the haptic feedback becomes. A key goal in the tactile layer is reducing the "noise" and producing the clearest haptic feedback signal.

The tactile layer is subtle. The tactile layer is no more than a few centimeters deep in physical space -- extending from the subdermal skin to the ends of our hair follicles -- but this depth limitation means that the entire tactile layer is jam-packed with information. Touch helps us acquire information regarding textures and shapes -- we can infer materials and identify objects in blind touch. [33], [34] Touch is central to interpersonal interactions, with touch playing a role in love, comfort, arousal, and the conveyance of emotion and sexual desire [33], [35]. Touch is bi-directional. These few cm are powerful enough to not only convey external desire, but to also produce an effect -- eliciting an emotional state than mirrors the input. Touch is a language into itself.[36] It changes the way we feel, and changes the way we perceive.

Touch is a danger sense. The evolutionary origin of touch is to avoid pain and damage to the user. We use touch to determine when our skin is breached, or when we're taking physical damage. Touch also tells us about potential dangers through the strong links to memory. We release rotted fruit instinctively, knowing somehow that it just "feels wrong".

When IVEs design for the Tactile Layer, they utilize the importance of touch. While the technology to sell the magic of visual and audio illusions is understood, touch experiences in VR are lacking the same level of fidelity. Without the ability to reliably reach out and touch the digital world, the experience of VR loses it is compelling nature. A successful IVE uses the tactile layer to augment the visual and auditory layers. It conveys energy and motion through vibration, and gives digital objects presence. The tactile layer in an IVE grounds us -- it gives the boundaries of the digital space physical consequence through impacts. The tactile layer drives our emotional state, creating fear, or desire. It makes us feel harried, or relaxed. Without the use of the tactile layer, the digital doesn't feel right, and full immersion becomes impossible.

LAYER 04: THE TASTE AND SMELL

The fourth layer, taste, and smell remain the unexploited Wild West of IVEs, with little to no technology-focused on addressing the fourth layer. Taste and smell are firmly linked, with our perception of food, a combination of these two sensory inputs. Tastants, chemicals in foods, are detected by taste buds, which consist of individual sensory cells around the mouth. Similarly, specialized cells in the nose pick up odorants, airborne odor molecules. Odorants stimulate receptor proteins found on hairlike cilia in the nose. Ultimately, in our brain taste and smell converge, creating flavor.

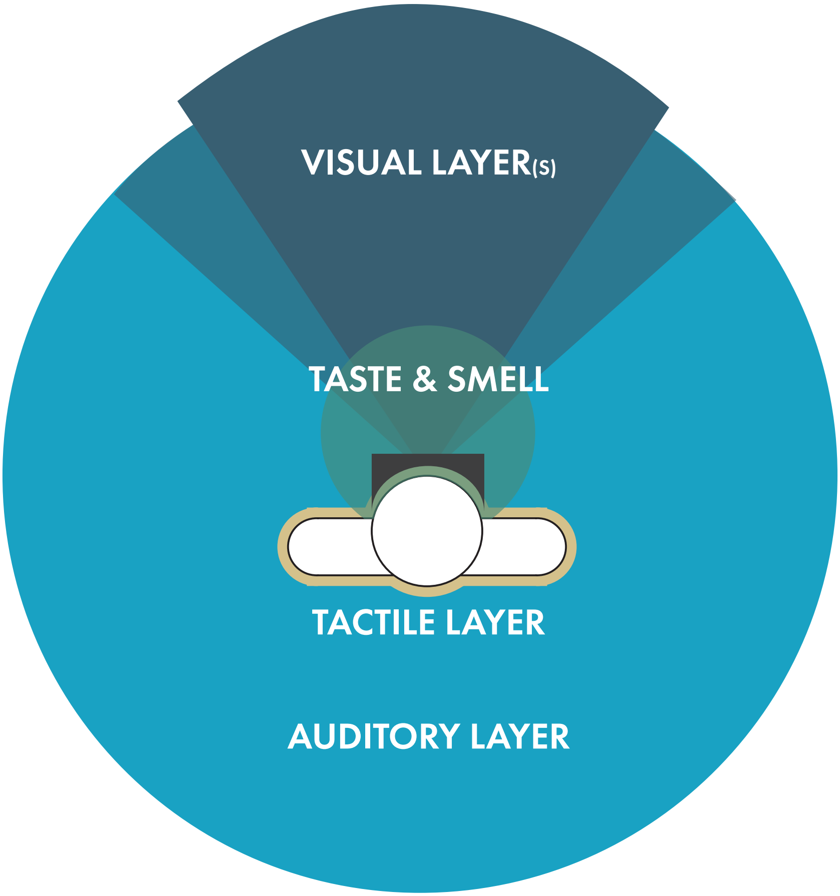

Fig: Layers of a compelling IVE superimposed.

The idea of smell-o-vision goes back to the early 20th century. In 1929, during the showing of The Broadway Melody, a New York City theater sprayed perfume from the ceiling. Arthur Mayer installed an in-theater smell system in Paramount's Rialto Theater on Broadway in 1933, which he used to deliver odors during a film. Smell-O-Rama, by GE, was used in 1953 to demonstrate a 3D rose with scent.

None of these early experimental smell layer experiments ever took off. In an IVE, smell, and taste play the smallest role in creating a compelling illusion. The use of smell and taste in an IVE is to provide context -- cueing a user to a emotional or environmental state.

Yet, there is potential in taste and smell layer augmentation. Both taste and smell have strong emotional ties and are strongly linked to memory. Smells can make us relax, or quickly disgusted. Imagine the depth of interaction when users can genuinely feel themselves on a racing track and smell burned rubber. Imagine being able to grasp the feeling of being on a battlefield complete with the intense gunpowder odor.

Very few IVEs design for taste and smell, but for a select group of use cases, this layer can make or break the experience.

CONCLUSION: THE INTERACTION CENTRIC APPROACH

THE LIMITATIONS OF THE INTERACTION CENTRIC APPROACH

An interaction centric approach is an approach to haptics built around the needs of IVE. While the body-centric approach works well for biology and psychophysics, and the tech-centric approach works well for engineering, the design, and creation of IVEs requires a new set of priorities.

The essential advantage of using the interaction centric approach is the outcome orientation. The interaction centric approach starts and ends with the user perception and experience. It explores what information is critical to a users understanding of an experience, and removes unnecessary noise from the haptic signal.

The stripped-down approach looks at what haptic cues are essential to the user's ability to understand the interaction. Superfluous information is discarded, allowing abstraction of haptic features instead. Rather than prioritizing a true to real-world recreation of every available stimulus, an outcome borne from both body-centric and tech-centric approaches, the interaction centric approach pulls from the Dieter Rams' ethos. "Good [haptic] design is as little as possible. Less, but better, because it concentrates on the essential [sensations], and the products are not burdened with non-essentials.

The IVE provides the designer with complete control over construct and sensory input. This is the core premise of the interaction centric approach. With full power, we can tailor the stimuli sent to every sensory channel. We can use haptics as an augmentation rather than a driver of the experience. The limits of the interaction centric approach become apparent in blind manipulation tasks -- where we cannot see the object.

> THE INTERACTION CENTRIC APPROACH STARTS AND ENDS WITH THE USER PERCEPTION AND EXPERIENCE. IT EXPLORES WHAT INFORMATION IS CRITICAL TO A USERS UNDERSTANDING OF AN EXPERIENCE, AND REMOVES UNNECESSARY NOISE FROM THE HAPTIC SIGNAL.

Without the visual cues, we cannot trade realism for immersion, and cannot exchange information from one layer to another. The interaction-centric approach relies on the more mature and advanced visual and auditory layers to drive fulfillment of the experiential objectives.

Without those layers, the results of an interaction centric approach are no different than the results of body-centric or tech-centric approaches. At the interaction centric approach should offer the ability to streamline haptic feedback to the most informative and essential elements. At it's worst, the interaction-centric approach offers parity with prior approaches.

THE INTERACTION CENTRIC APPROACH TO IVEs

With the knowledge of the limitations, the key question is how to implement the approach. The interaction-centric approach to IVEs is about deconstruction and reconstitution. We first break down the interactions we wish to accomplish into sub-parts. When applying to an IVE, rather than look at an interaction in isolation; we have to include contextual information that can help provide direction to the user journey and product experience.

Once we have a full understanding of the objective of the interaction -- from the experience level -- and an understanding of the perceptions we wish to create -- the perception level -- we can begin to assign information to sensory channels. With a clear understanding of what interaction layers -- visual, auditory, haptic, and smell/ taste -- will drive what parts of interaction, we can reconstruct the experience to be the most straightforward, most digestible version possible.

The ability to reduce the mental demand required to understand an environment is the driving factor behind immersion. The more brainpower we expend on interpreting a series of stimuli, the less brainpower we have devoted to interacting with the IVE. Haptic cuing values abstraction, asking not "does this feedback feel realistic" but "does this feedback improve immersion by intuitively guiding behavior.

THE INTERACTION CENTRIC PROCESS FLOW

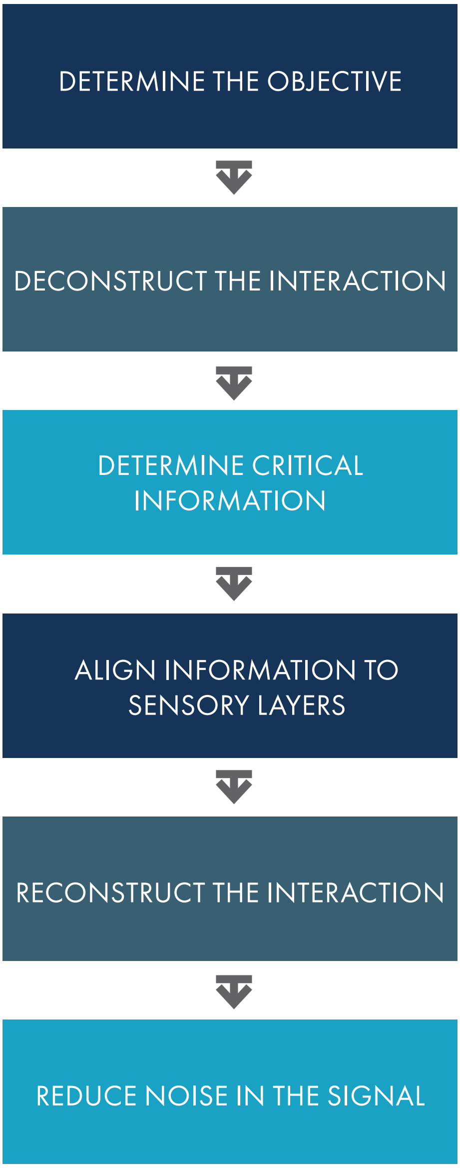

I use this framework again to outline how my approach to haptics works in the real world, and in the context of the Oculus inTouch controller.

HAPTIC SKETCHING -- DEVELOPING CONCEPTS THROUGH HARDWARE

CHAPTER 07

THREE CONCEPTS

My first visit to Oculus marked the completion of my research phase and the beginning of my prototyping phase. From the literature review and my prior work, I created three concepts of how to improve the haptic feedback within the Oculus Touch Controller. Each concept aimed to recover some of the interactions initially selected and would help drive my prototyping stage.

The first idea -- Button -- targeted the improvement of the controls category. This concept created a stronger link between the physical button and the vibrotactile feedback, putting a wideband VCA actuator inside the controller and responding to the position of grip and main trigger buttons.

The second idea -- Panel -- targeted the Gaming and Assembly categories. This concept tried to create a more compelling experience by reducing the moving mass driven by the vibration motor.

The third concept -- Zone -- targeted the Motions category. Using apparent haptic motion, two vibration motors could play a waveform/vibration slightly out of sync. This sync delay causes the sensation of something sliding through the hand.

I now needed to validate these concepts -- find some way or method of verifying these loose concepts, and embodying them into something tangible to iterate.

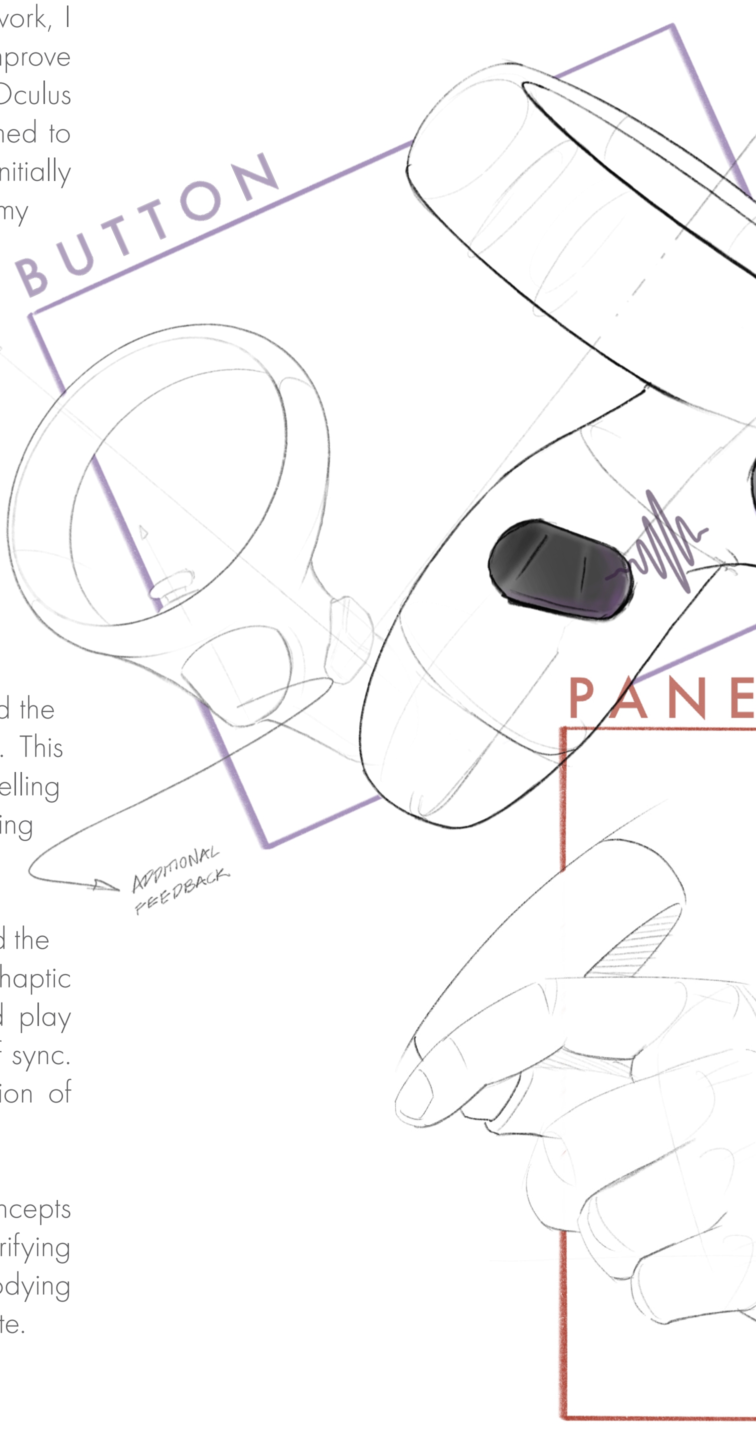

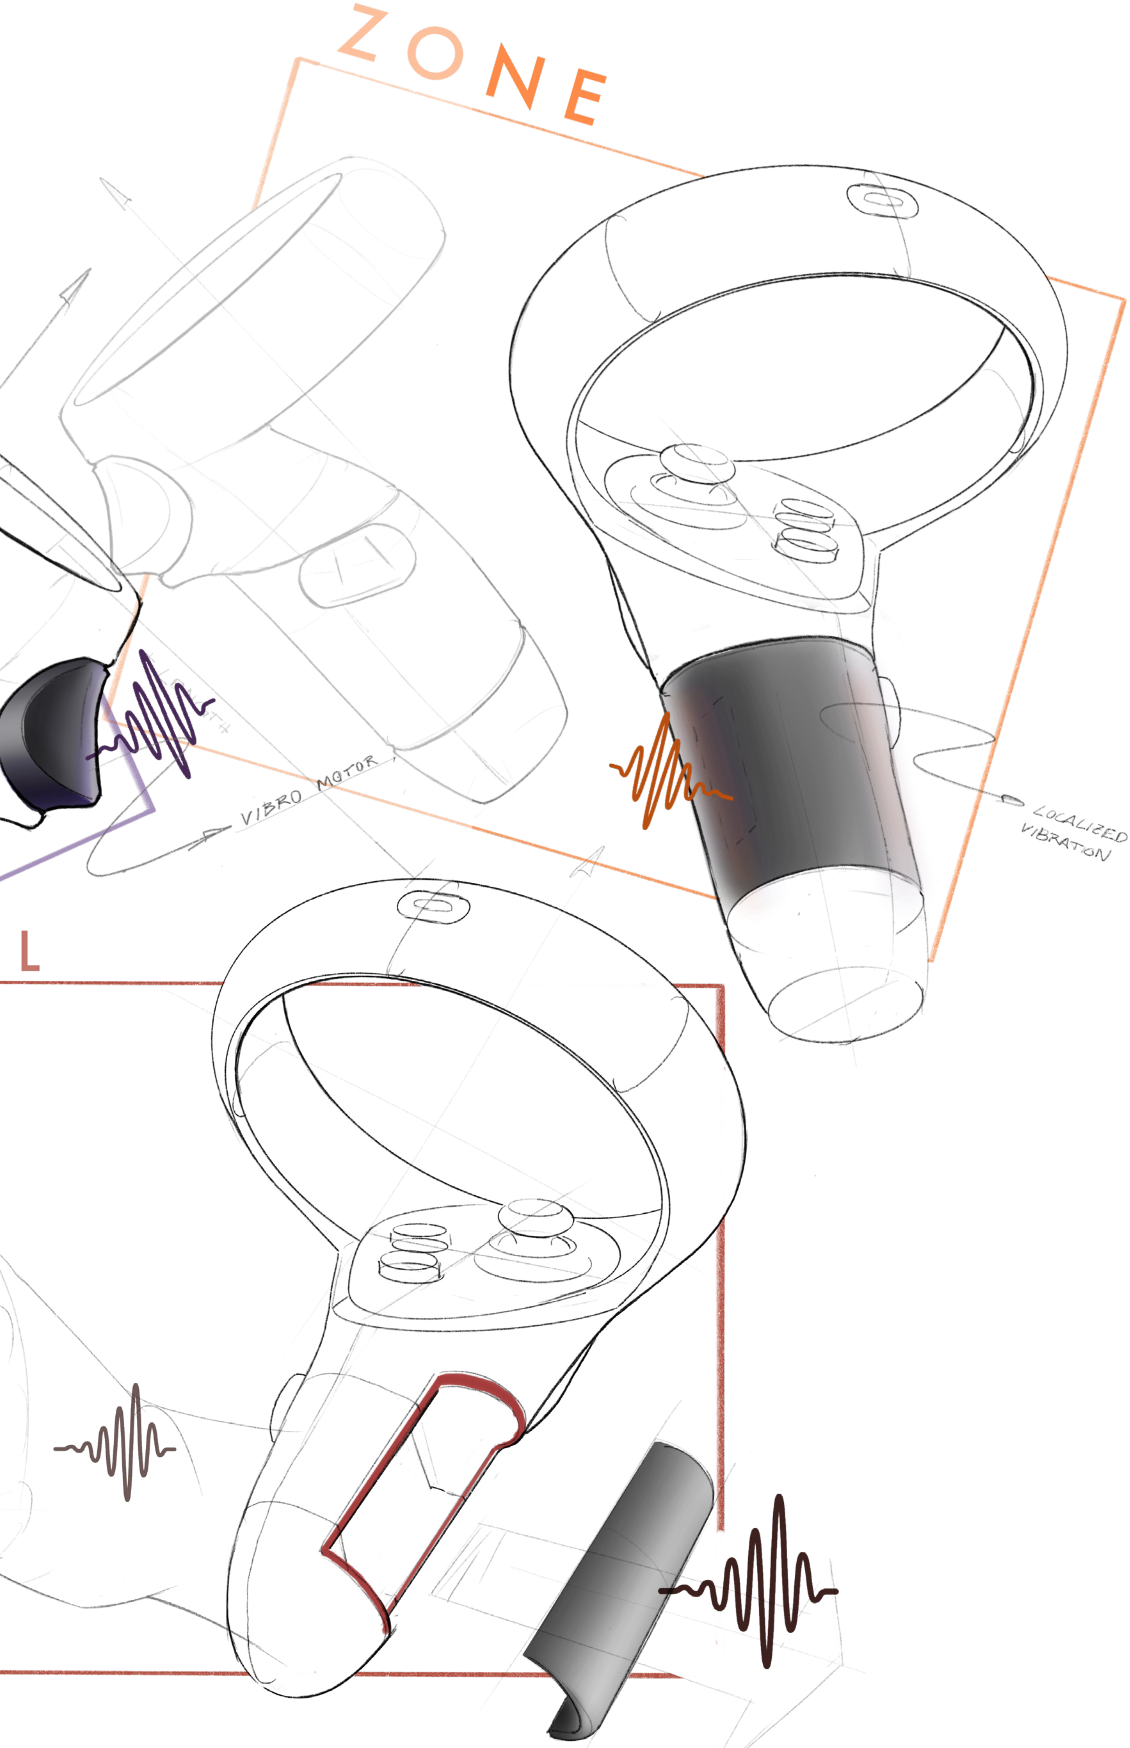
ZONE
VIBRO MOTOR
LOCALIZED VIBRATION
L

INTRODUCTION TO HAPTIC SKETCHING

Fig: Left: Simple Haptics, by Camille Moussette. Right: Examples of Haptic Sketches built by Camille with Microsoft Research.

Camille Moussette introduced the concept of a haptic sketch in his Ph.D. thesis -- Simple Haptics. [19] Just like how an industrial designer might use drawing and sketching to explore visual layer concepts such as product form, shape, and function, the haptic designer should have a similar tool -- something low cost, quickly iterable that allowed for the exploration of the tactile layer.

His dissertation outlined an approach to haptics that diverged from the clinical and scientific world of psychophysics and explored through rapid and physical prototypes the user experience of haptics.

"In [Camille's] developing view, haptic interaction design should strive to provide relevant, appropriate, and satisfying haptic stimuli that leverage materials, hardware, and software that are accessible and easy to craft and modify, leave room for design variations, and that aim to focus on purposefulness over technical prowess."

I resolved to use some of his principles in my thesis. I knew from my background with SenseGlove that things that worked "in theory" rarely fully translated into the real world. His "build to understand" approach appealed to my hardware background and prior work experience.

BUILDING A FOUNDATION - AN ARDUINO SHIELD FOR HAPTICS

One of Camille's key takeaways from his work at Microsoft Research was that a standardized set of I/O, controller boards, and connectors enabled prototyping. With a single interface code, parts, and lessons could be carried over from one sketch to another. Additionally, the standard interface allowed for I began to search for my standard layout. I set the following requirements:

The interface had to have either no-coding or Arduino/C++. It had been quite a few years that I had written any code other than some Arduino sketches, and I wanted to be able to quickly iterate without the need for support.

The interface had to be able to drivable without custom waveforms. While I had a method for quickly prototyping the sensations perceived during a haptic interaction, I didn't want to be limited to my recordings. This meant that Audio-to-Haptics was needed and that a PWM input was also desirable. The addition of built-in waveform libraries would also be a bonus.

The interface had to be low cost (<40€). I wanted to be able to build a few dozen haptic sketches, and not have to cannibalize one to build another. Each haptic sketch would have value and insights that I'd want to be able to revisit and compare to others quickly. If I only had a single controller/interface, it'd be hard to trial one after another.

The interface had to drive three or more vibration motors simultaneously. I wanted to be able to experiment with concepts such as apparent haptic motion, which required multiple vibrations played within milliseconds of one another.

With these requirements laid out, I began to search. The first requirement limited me to the (relatively large) Arduino ecosystem. Having already used the RedBoard Arduino Uno clone for the two interaction recorders , I decided to stick with the ATmega328 microprocessor and the Arduino Uno form-factor. Differences between the RedBoard and the standard UNO can be found at (https://learn.sparkfun.com/tutorials/redboard-vs-uno/all)

With the microprocessor chosen, I now knew how I was going to command the vibration motors, but not how I was going to drive them. Since I was planning on using LRAs for a lot of my sketches, I needed to use a specialty board that could provide the 1.8VAC they preferred (for details on different types of motors, see section "Types of Vibrotactile Actuators") I looked at the Texas Instruments DRV range, and the Dialog DA72XX range, two common low cost driver boards.

I settled on the DRV2605L, a driver with closed-loop feedback, audio-to-haptics, PWM input, and I2C control of a built-in haptic library -- the TouchSense 2200 from Immersion. I'd used the DRV260X series before at SenseGlove -- namely the DRV2603, for similar PWM only prototypes, and

BUILDING A FOUNDATION - AN ARDUINO SHIELD FOR HAPTICS (CONT'D)

the DRV2605L which we'd implemented in the SenseGlove Developer Kit 1.2. This familiarity and the ease of finding breakout-boards gave it the edge over Dialog's products.

When I looked closely at my requirements, I realized that a protoboard with multiple breakout boards wasn't going to meet my needs. I wanted something that was quickly swappable, compact, and lightweight -- I knew I'd be wearing some of my sketches, and a protoboard wasn't reliable for that use case. I began drawing out schematics for a custom PCB. It would use the Arduino shield concept -- a backpack that followed the UNO form factor and plugs in via the standard UNO I/O footprint. On the PCB I'd have an I2C multiplexer -- the I2C address for each DRV2605L is hard-coded and identical -- and four DRV2605L's. The DRV2605L would connect to the Arduino via I2C, and each has a PWM pin assigned. I also installed two TRS 3.5mm audio ports. Each 2605L would have a mono channel from the audio jack, so the two jacks could provide four unique signals at the same time.

This standard platform did limit the design space for my haptic sketches a bit, but the standard design's advantages outweighed that restriction. With support from the TU Delft Open Hardware initiative, I was able to design a custom PCB that would be shareable with others that were interested in building their own haptic sketches and could provide a tool for future use.

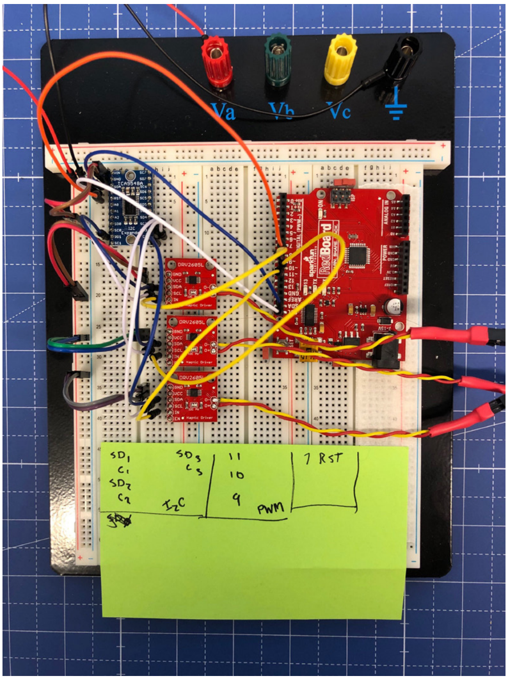

Fig: Prototype of Haptic Shield on a breadboard.

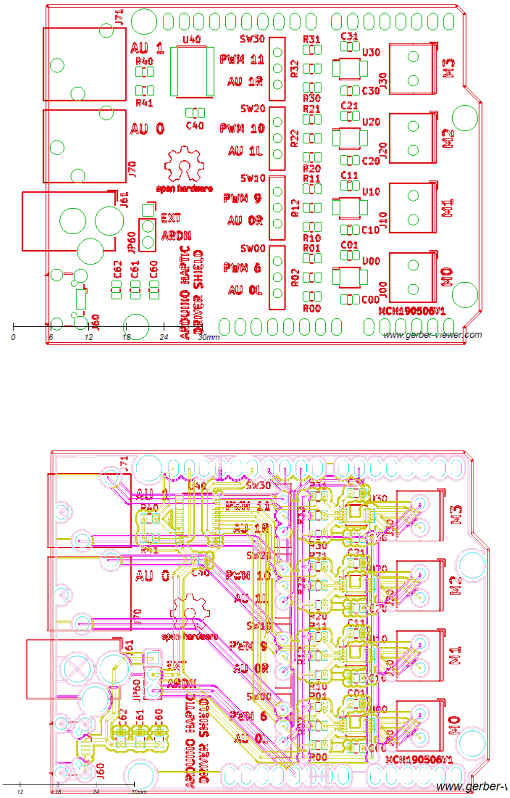

Fig: Top: PCB board layout of components for Haptic Shield
Bottom: PCB traces for Haptic Shield

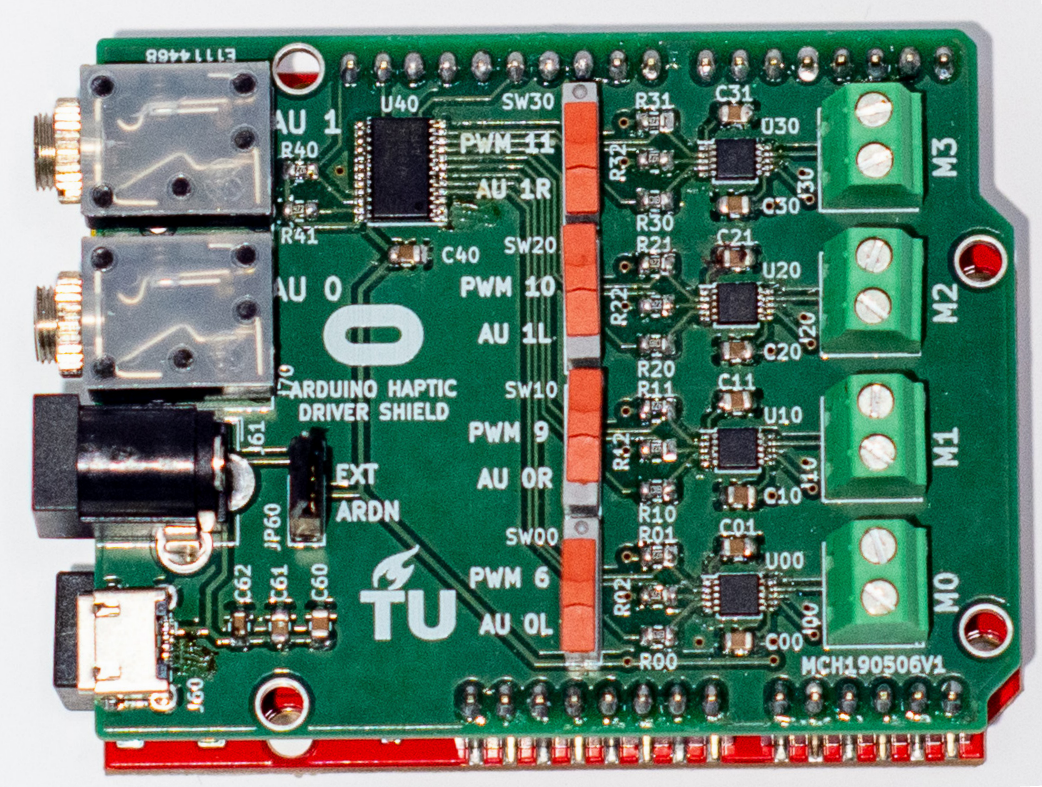

Fig: Top: Haptic Shield mounted on Arduino. Note switches for audio jack input or PWM input
Bottom: Haptic Shield with schematic of the PCB.

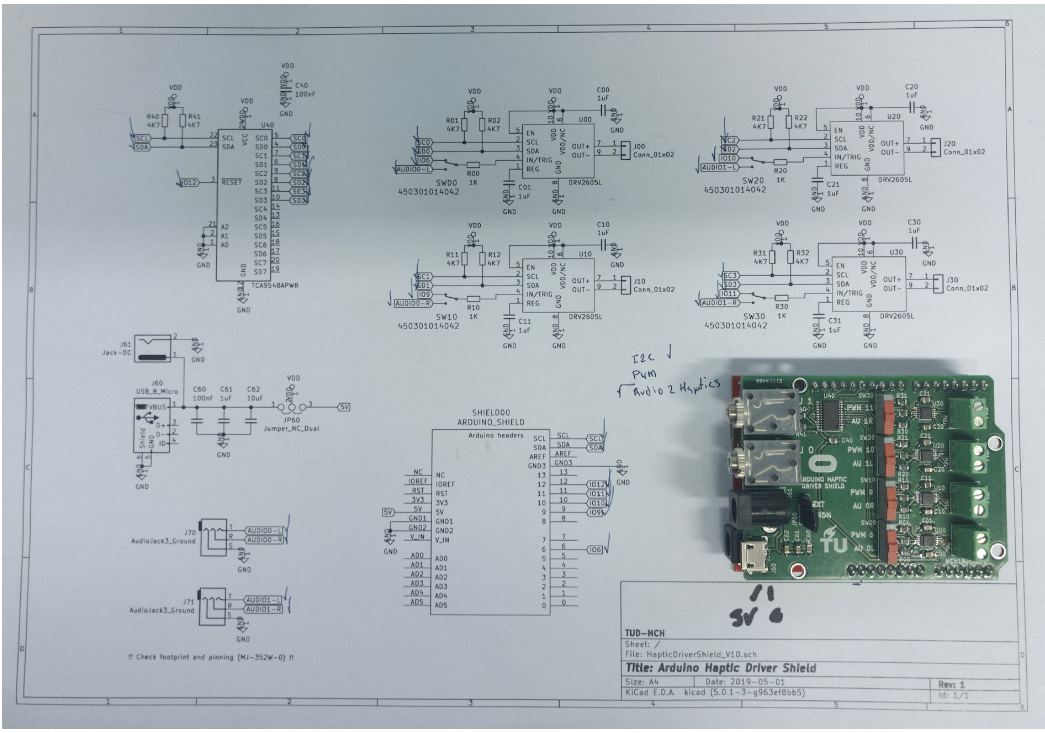

SKETCH GROUP 1: CONTROLLER SKETCHES (BAMBI)

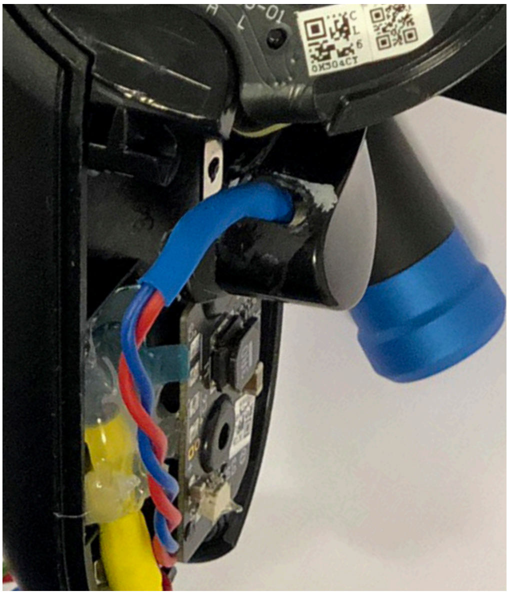

Fig: LRA installed inside Oculus Touch grip trigger

Fig: LRA installed inside Oculus Touch front trigger

My first set of sketches was designed to validate some of the concepts discussed during my Oculus visit. While the concepts were to take advantage of more advanced haptics/vibration motors, I felt that the basics could be explored with Simple LRAs. This haptic sketch group was installed on a disassembled Oculus Touch controller, used the Jin Long LVM Series LV081230B-L30 LRA or Jin Long G0832013D LRA for effects, and the Haptics Shield for control. This series was named after characters in the Disney story Bambi in honor of "Thumper", the infamous bunny known for pounding the ground with his foot.

SKETCH 01 - BUTTON AUGMENTATION (FALINE)

The first sketch, Faline, explored ideas outlined by Ogawa et. al, and Shim et, al -- namely the ability to change the perceived physical features of a button through vibrotactile augmentation. The Oculus Touch controller trigger and grip buttons were disassembled, and a G0832013D was installed. The button's plastic shells required slight modification to run power cabling to the motors but were otherwise unchanged. The LRA's were mounted

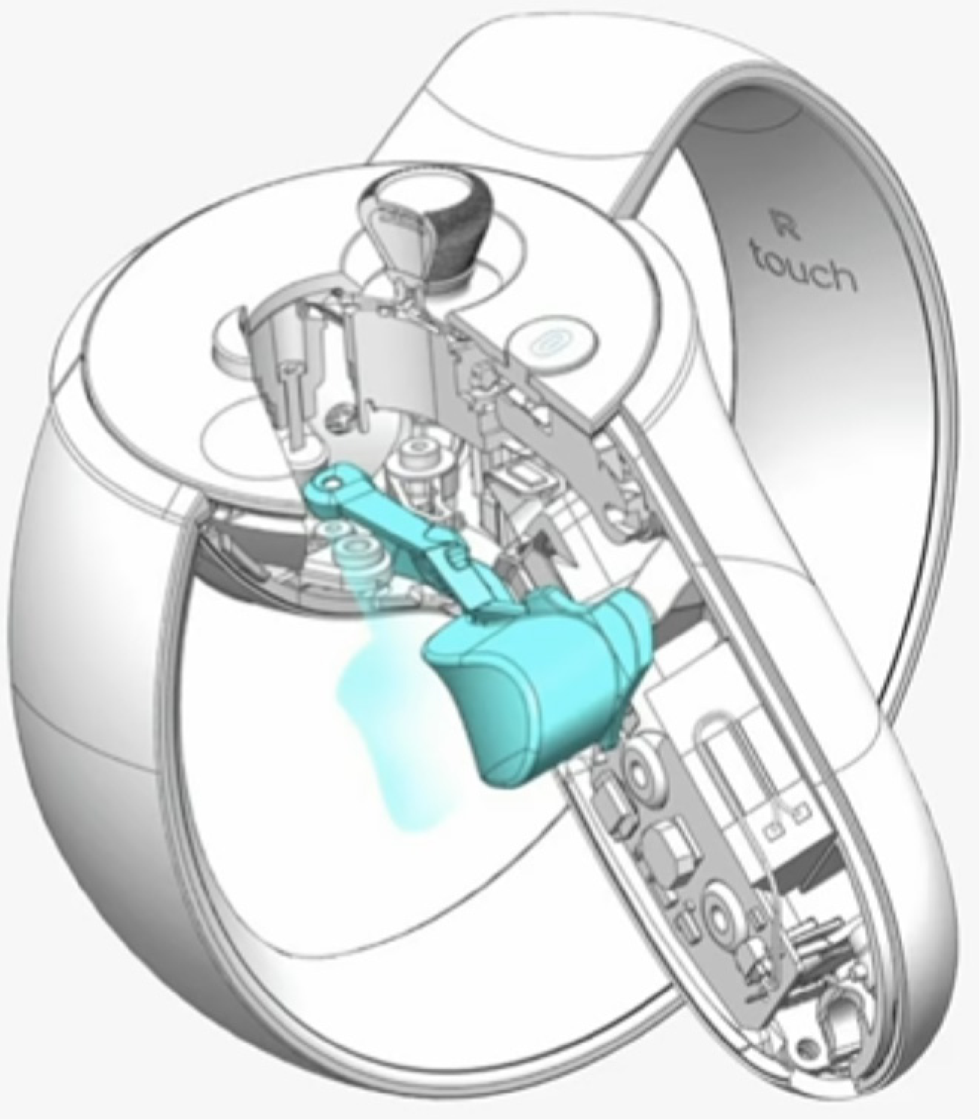

Fig: Oculus Touch cutaway showing existing 2DLRA and grip trigger (teal)

perpendicular to the button, in order to provide vibrations directly into the finger.

I had two main goals in mind for this test. First, I wanted to see if the grip or trigger buttons would feel different if I played a vibration during use. The second was the see if I could wholly replace any of the buttons with alternatives -- freeing up valuable internal space.

This sketch was quite simple, but provided clear, compelling sensations. As isolated assemblies, the buttons are naturally isolated from the rest of the controller. Therefore, vibrations within the buttons felt clear, crisp and compelling. The playback of clicks and thumbs from the DRV2605L's built in haptic library (Touch Sense 2200) could be distinguished from other vibrations inside the controller. While the ability to change texture wasn't fully implemented, light vibrations while gripping the controller gave me a perspective of what could be possible.

Fig: Sketch 01-03 on Oculus Touch controller.
Blue/Red wires - Sketch 01
Green/Red wires - Sketch 02
Orange/Red wires - Sketch 03

SKETCH 02 - PANEL ISOLATION (THUMPER)

The second sketch, Faline, explored material properties and vibration isolation to attempt to create a bit more thump from a vibration motor. In this sketch, the Oculus Touch controller's battery compartment was removed and modified with a silicone layer that was soft and deformable. With a hardness of Shore 65A, this dampening layer allowed the battery cover to vibrate independently of the rest of the controller. On the back of the battery case, I affixed a LV081230B-L30 LRA with superglue (cyanoacrylate adhesive).

The hope was that by reducing the mass being moved by a vibration motor we could increase the perceived effectiveness. Most vibration motors are tested with a reference mass of 100g. If I could reduce the mass that the vibration motor had to actuate, then I could create a more targeted and higher impact. The goal was to only vibrate the minimum amount of mass that was in contact with the palm. But removing the battery door and allowing it to vibrate independently, I reduced the affected mass from ~160g to ~ 10g

Overall this test was successful. I could feel a distinct difference between the two vibration motors. Combined with the lessons learned from Faline, I knew that isolation of assemblies and panels would be worth pursuing.

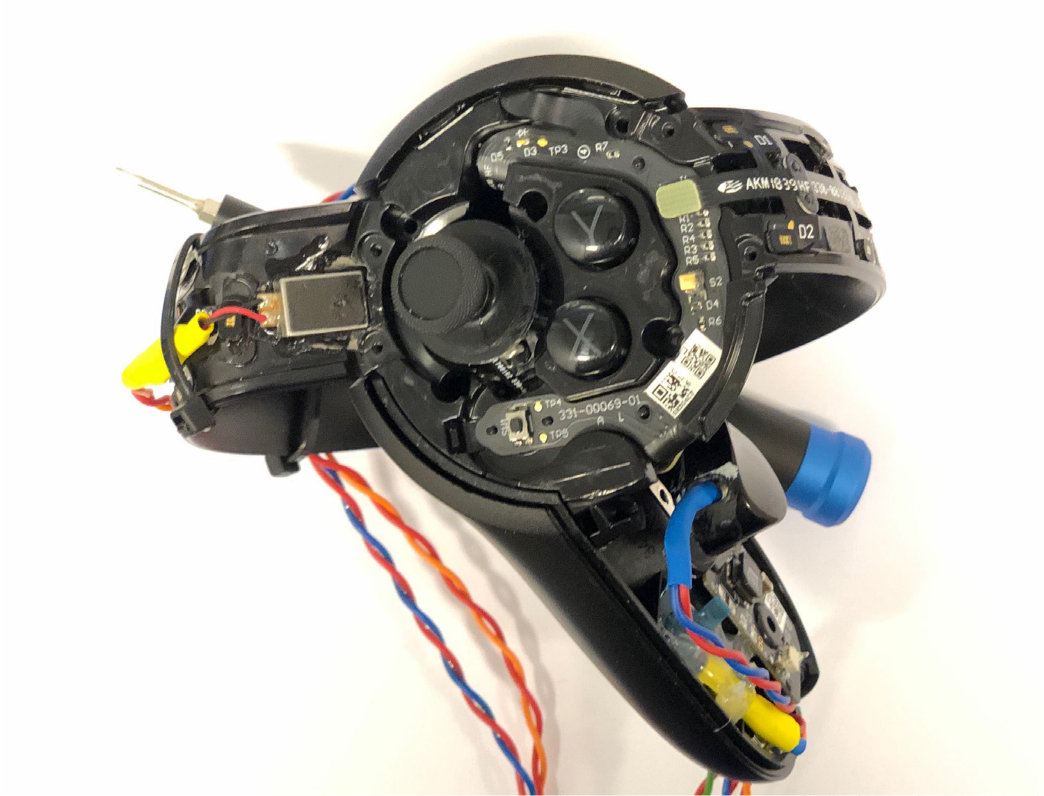

Fig: Apparent haptic motion on Oculus Touch. LRA located next to joystick

SKETCH 03 - APPARENT HAPTIC MOTION (FLOWER)

The third sketch, Flower, explored a phenomenon known as apparent haptic motion. This neat haptic illusion creates a sensation of motion by using two vibration motors with similar intensities and matched frequencies played slightly (<70ms) out of sync. I affixed a LV081230B-L30 LRA with hot glue to the inner base of the grip, and another to the point where the grip met the LED ring for the constellation tracking. This maximized the distance between the two LRAs minimized interference and would allow for the largest range of timing variations.

I wanted to see if this was still possible when gripping and using the Oculus controller to distinguish the two pulses, and what useful ranges in time, intensity and frequency were possible. The hope was that through this sketch to get a series of variables and limits that could be converted into the sensation of stick/slip through the hand.

From this sketch I found that this effect was possible in a single mono-body style controller, but that isolation would likely augment the working ranges for the controller. I could perceive an effect between 40ms and 54ms. Within this range direction cues (up/down vs down/up) were clear, and it felt like one single vibration rather than two distinct interactions. While there are many more variables to explore, Bambi validated the idea of apparent haptic motion through a controller.

SKETCH GROUP 2: FREEHAND SKETCHES (LION KING)

My second set of sketches focused on vibration motor and position. While most literature suggests that position has a small impact on the imparted sensation, I wanted to validate and verify this through some haptic sketching. I also wanted to validate some literature that suggested the hand could be broken down into Hand and Finger zones of influence. These sketches aimed to perfect a single interaction -- the button click -- that Oculus had suggested during my first visit. All of these sketches used 3D mounts and the LoFelt L5 vibration motor.

The button waveform was a physics-based waveform generated from 3D acceleration data of an OMRON SPST I/O button. It was played through the LoFelt Evaluation Kit (EVK) with the Digital Signal Processing Engine (DSP) in Haptics Bypass Mode, meaning the waveform was not at all modified.

SKETCH 04 -FINGER POSITION (TIMON)

The first sketch in this series aimed to determine the sweet spot for vibration perception. One of the goals of this thesis is to produce a controller that allows users to interact with virtual buttons and control panels in an immersive and convincing manner.

When we interact with buttons, we use our fingertips. I wanted to see how far away from the finger I could present the vibration feedback of a button interaction and still have it "feel" satisfying, realistic, and immersive.

Users had a LoFelt L5 affixed to the hand on one of seven (7) positions (denoted by the bone below the skin)

› Index Distal Phalange,
› Index Intermediate Phalange
› Index Proximal Phalange,
› Metacarpal (Distal)
› Trapezoid / Metacarpal (Proximal End)
› Scaphoid
› Radius

They were then asked to "click" a virtual button in front of them with their eyes closed and report the satisfaction of the click, the realism of the click, and the immersive-ness of the experience.

After a quick study, it became clear that vibrations quickly deteriorated in all three factors after crossing the finger webbing. The sweet spot for vibrations to be associated with the fingertip spanned from positions 1-4. Interestingly, the wrist became another possible spot for vibrations. While not as good as the "sweet spot", users definitely expressed improved sensations. This was further explored in Sketch 07 - Zazu.

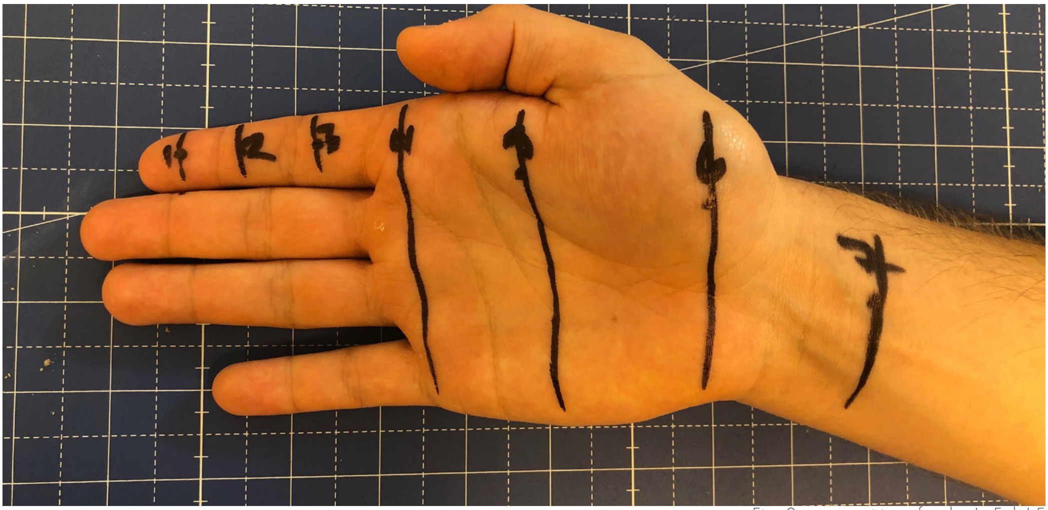

Fig: Seven positions for the LoFelt L5

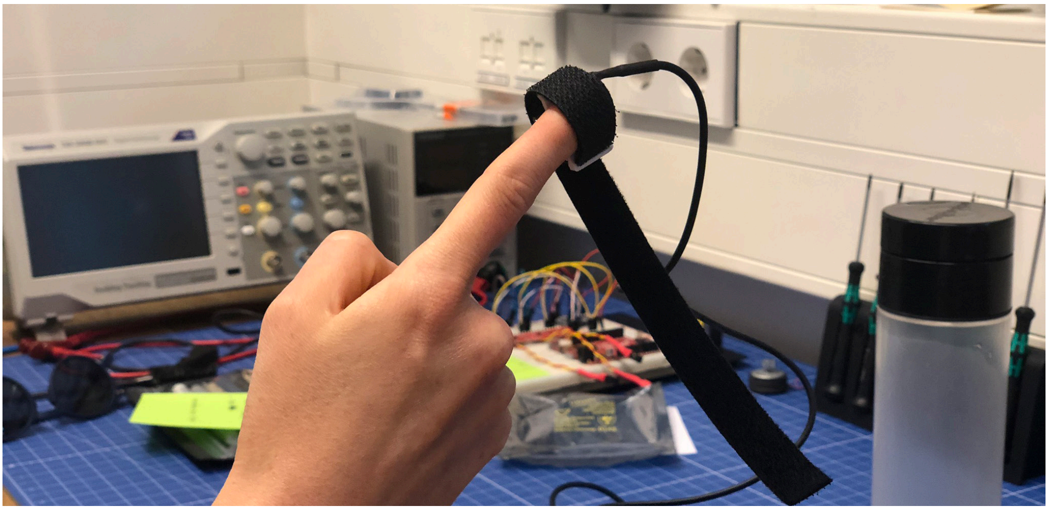

Fig: Reserach in progress with L5 on located in Position 1

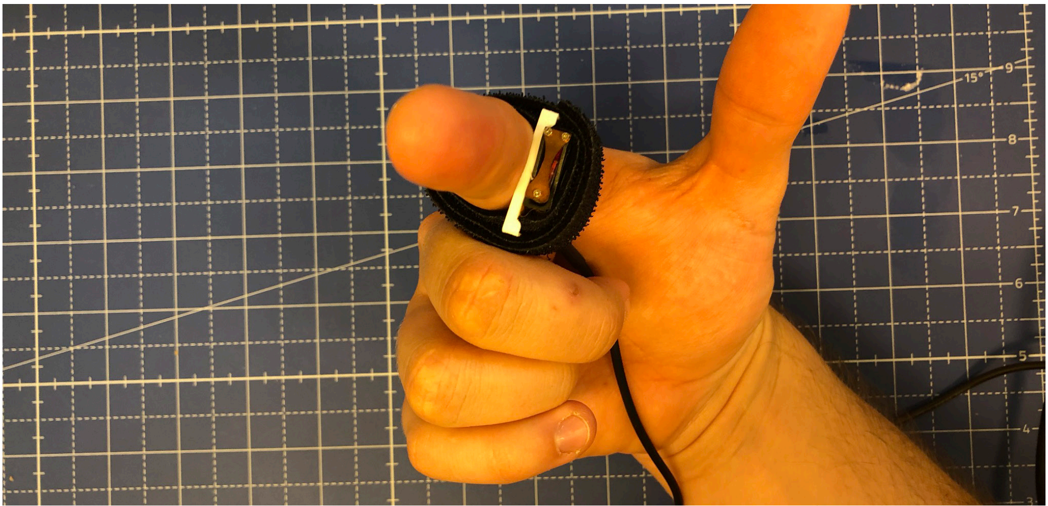

Fig: Reserach in progress with L5 on located in Position 3

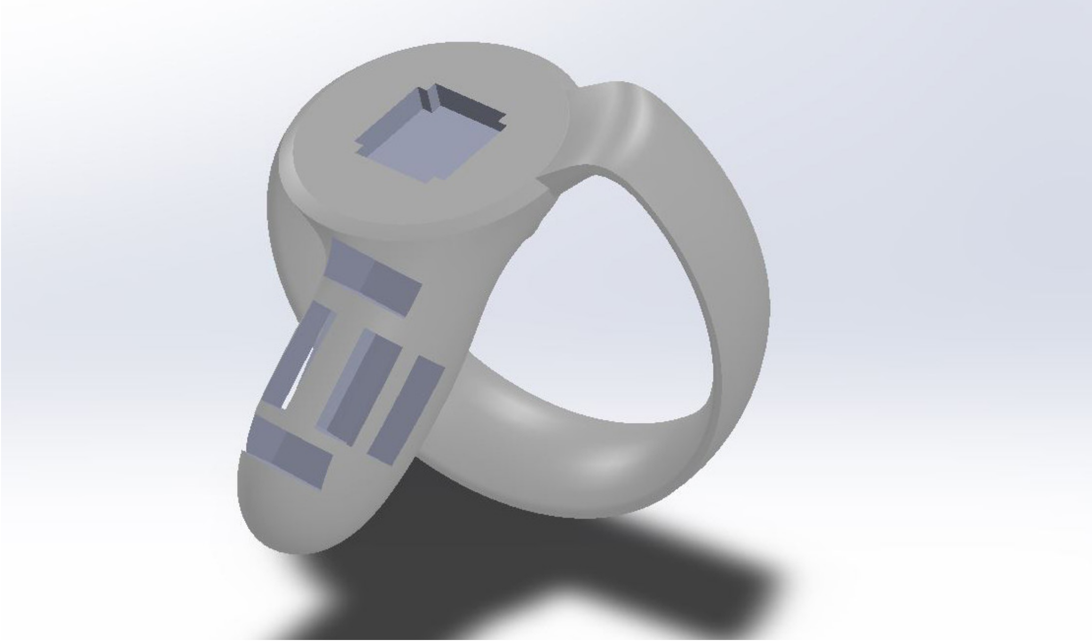

Fig: CAD Model of Sketch 05. Each slot sized for a LoFelt L5

SKETCH 05 - ROTATION POSITION (PUMBA)

In the second sketch of the series, I wanted to verify that the target force for the vibration motor (>3G) was sufficiently high that position and orientation of the vibration motor would have a minimal impact on the perceived intensity of the vibration.

A controller form was printed with an Ultimaker S5 FDM/FFF 3D printer. A series of slots were cut halfway down the grip, each 30 degrees offset from the prior. Additionally four slots were cut vertically along the grip axis of the controller. Each slot was sized to fit a LoFelt L5 wideband VCA. In total, Pumba contained eight unique positions or orientations for the LoFelt actuator.

To test the universality of the vibration effect, I played three Youtube videos -- one that contained a series of gunshots from Call of Duty WWII, a second that contained driving footage from Gran Turismo Sport, and a final one that contained sixteen unique buttons click sounds. The LoFelt EVK was used to turn the sound data from the videos into haptic interactions. I used the Gaming DSP mode, which had the most aggressive filtration to produce a more defined set of experiential highs and lows for each video. Each clip was played with the LoFelt L5 actuator in a unique position.

While I felt every position, more centralized locations in the grip felt better the extremes. I was expecting more consistent performance across the hand, given the strength of the LoFelt actuator; however, the wave propagation took a more significant role than expected. For an IVE, this would have to be taken into account. Actuators too far up or too low in the handle will cause some discongruity with the desired experience.

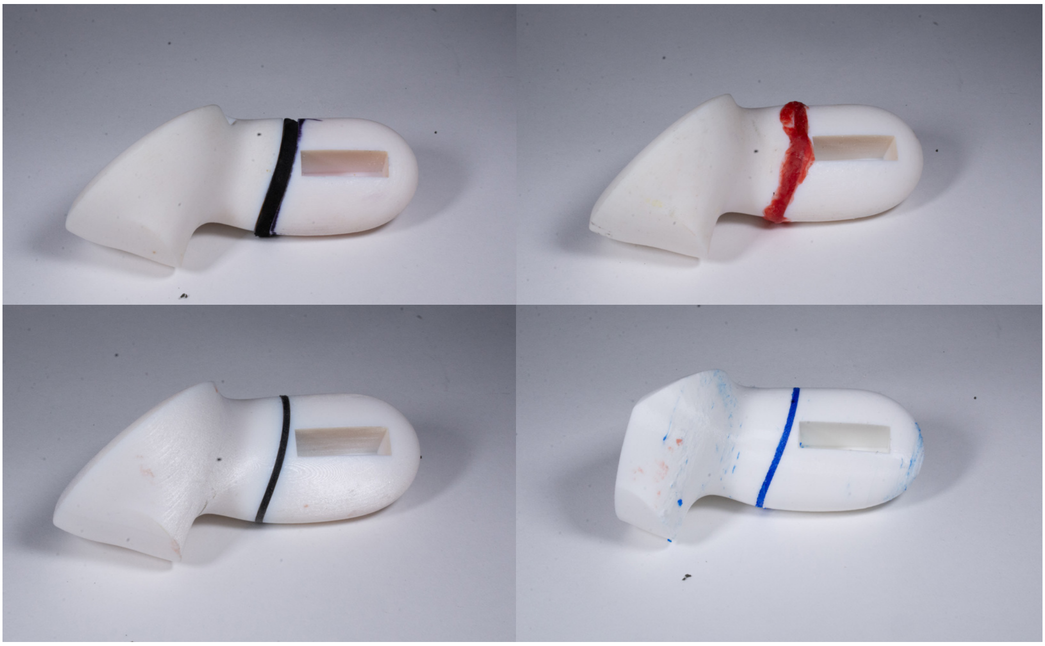
Fig: Sketch 06 in various iterations. Clockwise from top left: closed cell foam, cheese wax, TPU, Silicone.

SKETCH 06 - TPU ISOLATION (SCAR)

The third sketch in the series tried to combine ideas from Bambi and Thumper. I wanted to see if it was possible to isolate areas of the controller from one another through the use of soft elastomers. The controller was cut in half, with each section being made from a hard plastic material. A softer dampening material was put in between each half. The hope was that in splitting the controller two distinct signals could be presented to the upper and lower halves of the hand.

The first version of Scar was printed with an Ultimaker S5 FDM/FFF 3D printer. The boundary between the two hard plastics was made from Thermoplastic Polyurethane (TPU) with a Shore durometer hardness of 95A. By printing the boundary at 80% infill with a Gyroid pattern, the apparent hardness was lower -- closer to a Shore 85A. While this version was clearer to understand than the original Thumper sketch, there was still a significant amount of dampening between the two channels, and the apparent haptic motion effect wasn't compelling enough.

Three additional versions were built, using softer elastomers and thicker gaps to further dampen the two signals. A Shore 60A elastomer, cheese wax, and Shore 50A silicone all were used, but in the end, the signals always felt muddy.

Scar showed that there's potential in the concept of an apparent haptic motion, but given the limited range of motion and use, along with the manufacturing and material challenges, it simply wasn't viable for the limited time and larger scope of this project.

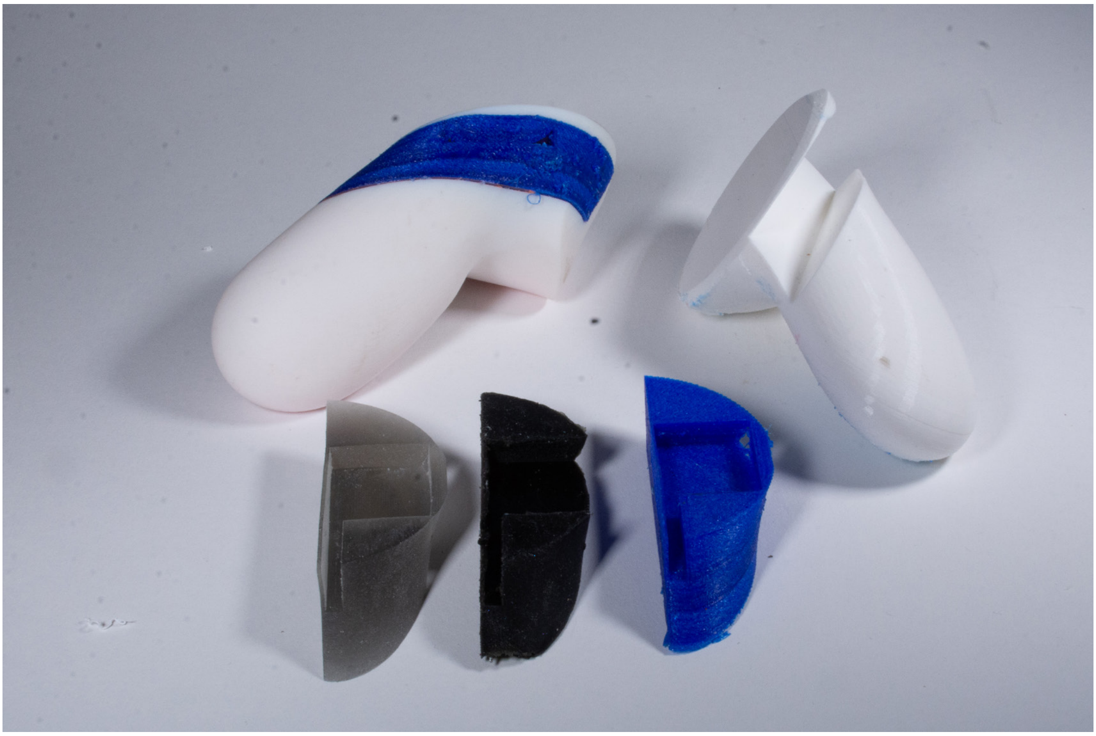

Fig: Sketch 07 in various iterations. Grey: Power Bed Sintered Material, Black: Cast Silicone, Blue TPU. Slot sized for LoFelt L5

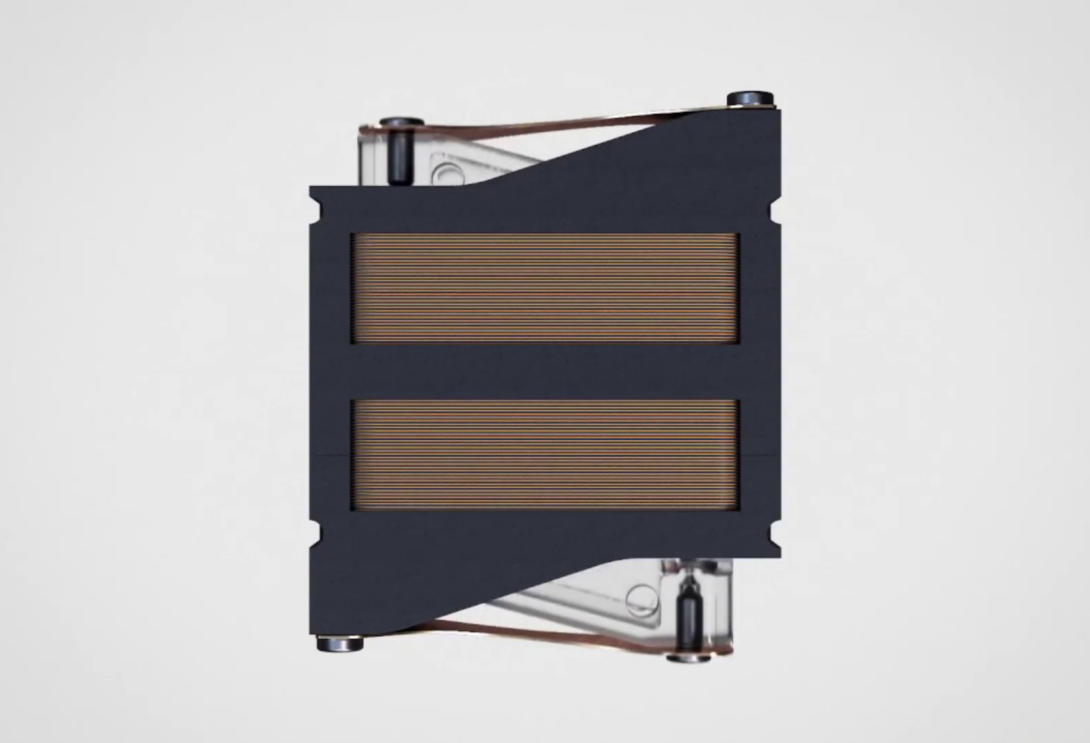

Fig: Render of the LoFelt L5 Actuator used in testing.

Fig: Sketch 07 in each iteration.

SKETCH 07 - CONTACT PATCH (ZAZU)

In sketch seven hoped to create a strong clean sensation of a button click down the hand. Prior work at SenseGlove led me to conclude that if we could create a contact area for the index finger with a material with similar vibration properties to the skin, it would create a clearer more compelling vibration. Ideally, this signal would travel down the length of the finger and create a localized sensation at the finger tip, and with that a more compelling button click.

The first version of Zazu was printed on the Ultimaker S5 with the same material properties as Sketch 06 - Scar. With a fairly stringy/rough surface finish, this print did provide some interesting effects due to the higher contact area between the hand and the controller, but vibrations were muddy and damped.

With the first iteration not displaying any of the desired effects, I asked the Oculus team to use some of their more advanced prototyping tools to try and find a material that had the appropriate material properties. The second iteration was printed on a HP 3D powder bed printer. This created a nice isotropic (same material behaviors in all directions) material, and was softer than the TPU. With a Durometer hardness of roughly Shore 60A, we were beginning to get into the right ballpark.

Sadly the second iteration was also underwhelming, representing a step back from the first. Without the textured surface, vibrations were even more damped and less satisfying. A final version was built by the Oculus team in Menlo Park using a cast silicon.

Overall, this sketch didn't produce the desired effect. While an interesting idea of trying to create a conformal contact patch, It simply didn't produce any compelling interaction or improvement.

SKETCH GROUP 3: JOYSTICK SKETCHES (ALADDIN)

My third set of sketches explored force and texture modification of a joystick through vibration. Originally presented by Smin and Lee [13], who used a piezoelectric actuators to vibrate the tip of a joystick, created an experience of higher or lower force simply by vibrating the joystick. I wanted to validate if this same effect could be achieved with traditional vibration motors -- LRAs -- and that the effect still worked even with the vibration moved away from the joystick tip.

The sketches were driven by a custom designed PCB I had designed during my time at SenseGlove, and was allowed to re-use for this project. The PCB was simple, containing some header pins, and a DRV 2603 -- a version of the TI driver family that was optimized for PWM driving, but could also have been done with the haptic backpack I'd developed. The DRV chip powered a Jin Long LVM Series LV081230B-L30 LRA.

SKETCH 08 - JOYSTICK PWM (PRINCE ALI)

The eighth sketch, Prince Ali explored the inter-relationship between joystick displacement and increasing vibration amplitude. The joystick's X and Y axis potentiometers were wired up as voltage dividers -- so that voltage output increased as joystick displacement increased. Each output was then connected to a analog pin on an arduino, allowing me to interpret distance from neutral on a scale from 0-1023. These two inputs were then used to create a linear from the neutral position, which scaled a PWM signal (0-255) as a function of distance from origin. The PWM signal then drove the frequency between vibrations -- the further from the origin, the more frequent the vibrations.

In effect, what this sketch hoped to understand was could the abstraction of a footfall be linked to joystick position. Each vibration represented a hypothetical footstep of an imaginary character who's movement was controlled by the joystick. The further from the origin, the faster the character "ran".

From this sketch I found that this effect was possible. While the abstract nature of the "footfall" did inhibit the compelling-ness of the effect, it was a very clear cue when presented to the hand. I knew something was going on, but it wasn't immediately clear what that was. Since the haptic sketch was an abstraction, I felt this limitation was able to be overcome with better waveform design, something more akin to a footstep than a simple buzz. This effect was also implementable immediately with no major hardware changes to the Oculus Quest Touch controller.

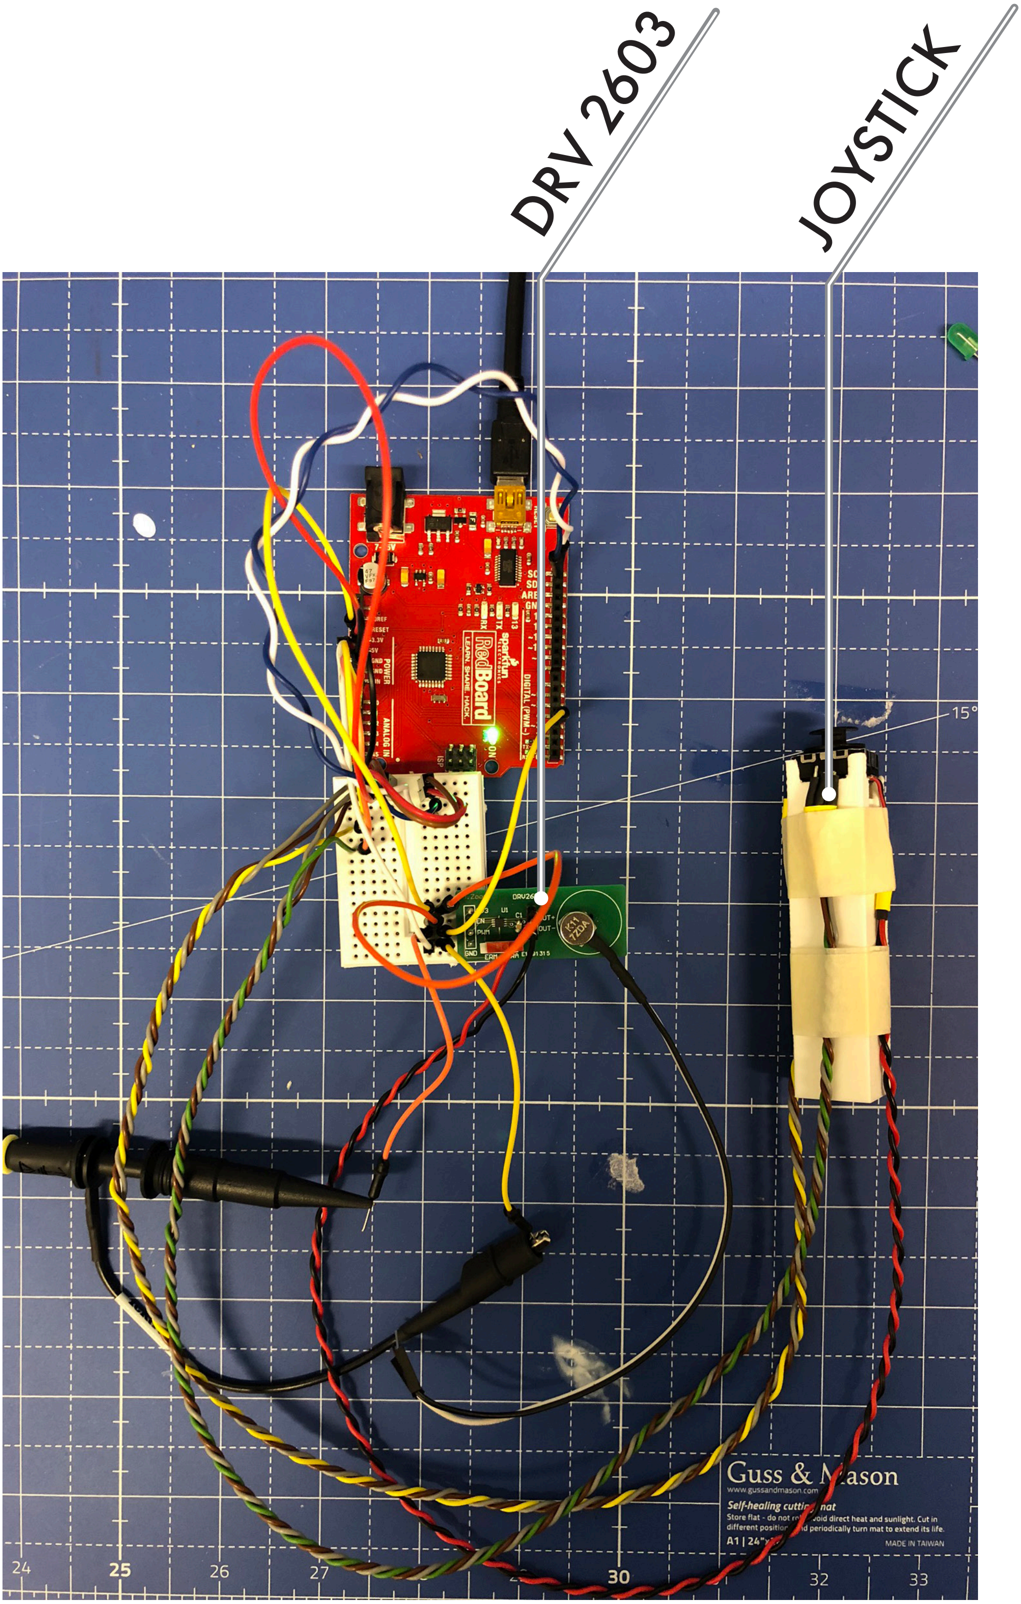

Fig: Sketch 08. PCB (Green) contains a DRV 2603. LRA is connected to Red/Black Wire.

SKETCH 09 - JOYSTICK FORCE FEEDBACK (GENIE)

The ninth sketch came about due through an accident. In rushing to write the arduino code for Sketch 08 -- Prince Ali, I had forgotten to include a two key lines of code:

```
drv.useLRA();
drv.selectLibrary(6);
```

These two lines told the haptic driver to drive the connected motor with 1.8V AC power (3.6 DC running as 1.8V and -1.8V) for LRAs rather than the 3.3V DC for ERMs. By wildly mistreating the LRA, I was producing high frequency vibrations way outside the normal operating range of the LRA. Interestingly, this high frequency noise felt like friction inside the joystick mechanism. The harder I pushed on the joystick, the more resistive the joystick felt.

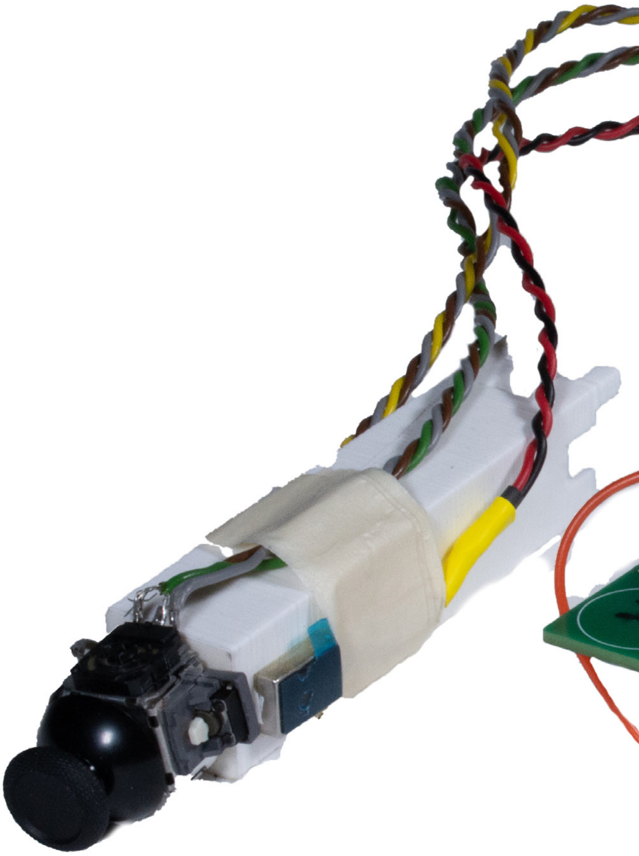

This happy little accident was compelling enough that I felt if I could replicate it with a wideband actuator that was able to comfortably operate at these sorts of frequencies, it was very much worth including into my final design.

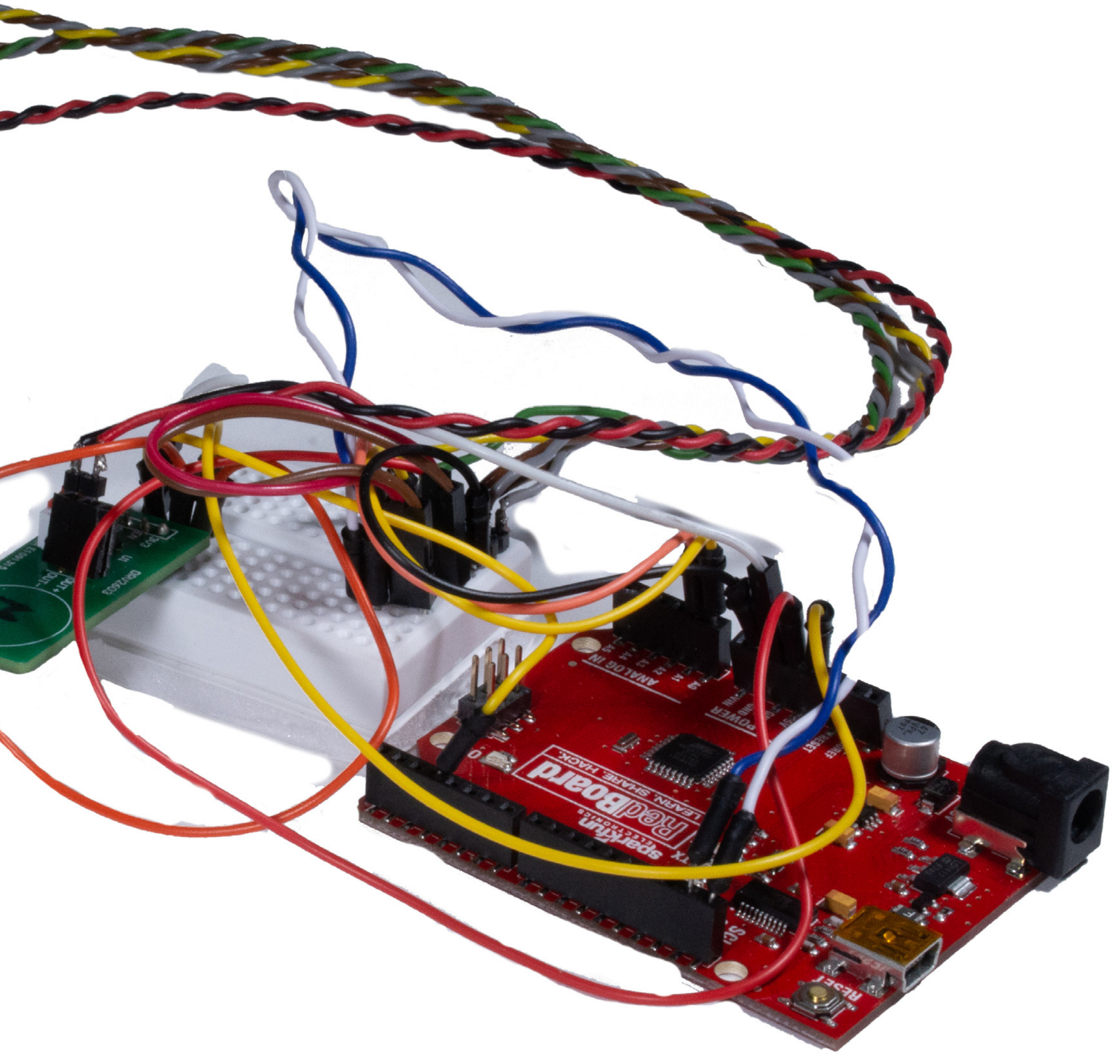

Fig: Haptic Sketch 08/09. Note LRA mounted on side of the plastic grip.
LRA position also tested on bottom , directly below joystick.

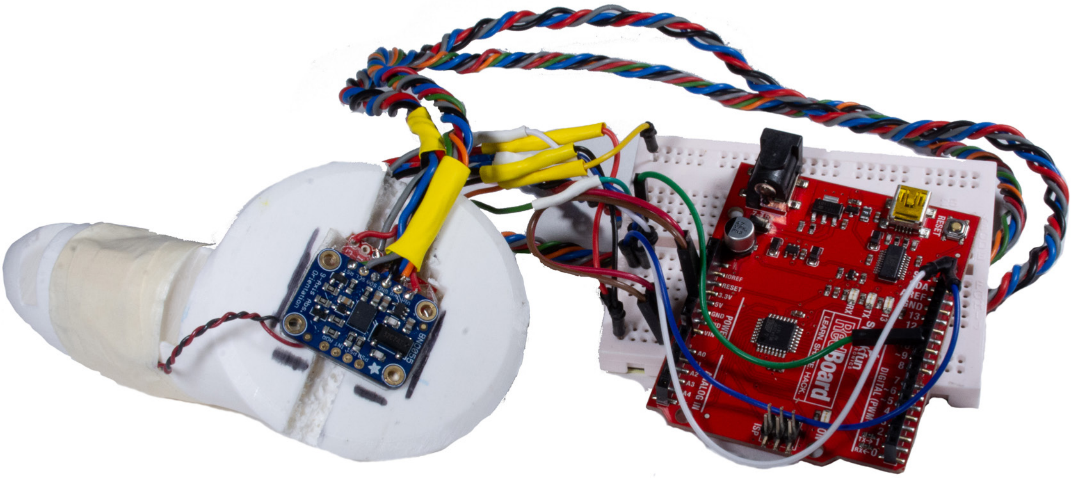

Fig: Haptic Sketch 10, with 2DLRA mounted inside grip handle.

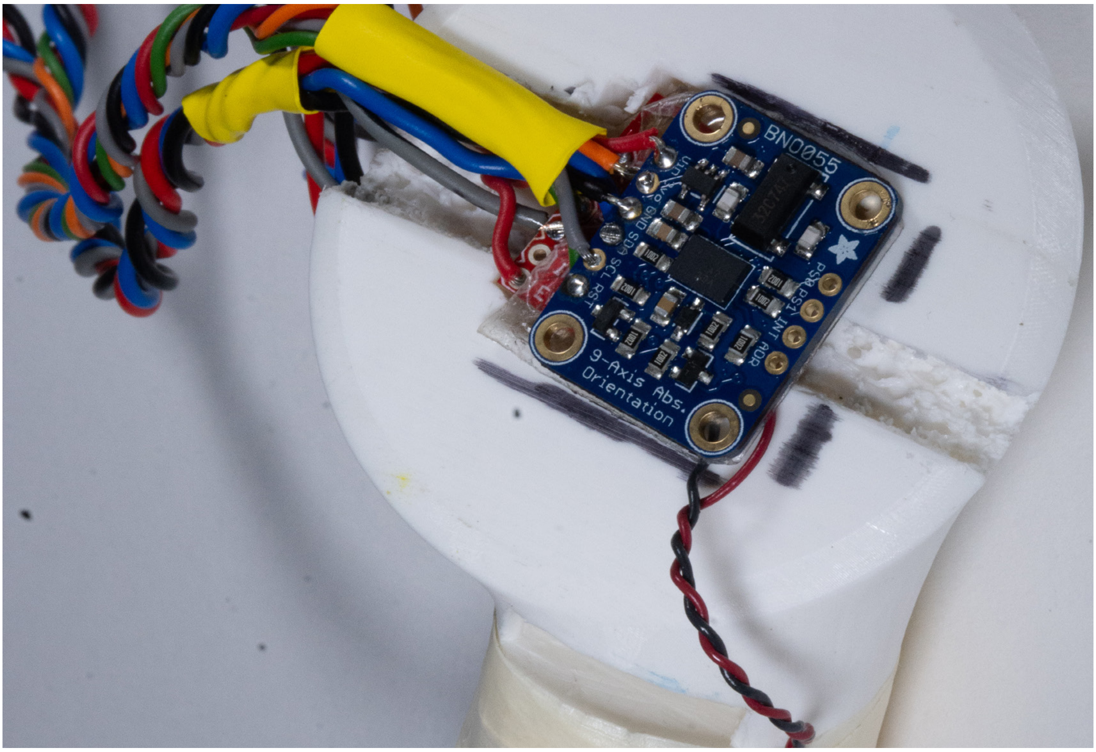

Fig: BNO055 IMU (BLUE) mounted above DRV2605 breakout board (RED)

SKETCH GROUP 4: IMU SKETCHES (MULAN)

The final sketch group explored how to create mid-air haptic experiences. The core of this set of sketches was the BNO055 Inertial Measurement Unit (IMU), which tracked acceleration, orientation, and rotation position in space. This tracking was relative, meaning that I could not place the hand in 3D space, but I could measure rough deviation from the preceding position. The IMU fed these data to an Arduino and vibrations were provided by the ALPS/ALPINE haptic reactor, a 2DLRA in the original Oculus Touch, HTC Vive Controller and HTC Vive Focus Pro Controller.

SKETCH 10 - CLICK-CLICK & IMU (MUSHU)

The final sketch I built was to test if it was possible to create a compelling in-air-haptic experience of a button click. The goal was to build a controller replica that allowed the user to make a clicking motion in mid air and evaluate how compelling the haptic only experience felt.

The IMU looked for an acceleration greater than 2 m/s in either X or Y orientations. Once detected, the arduino sent a command to a DRV 2605L driver to play a "click" waveform. The "click" alternated between a 100% "CLICK" and a 30% "click". The goal was to recreate the double click experience of a toggle button.

The effect worked at a small range of velocities, but was quite compelling at the right speeds and ranges. The effect was further improved with eyes closed and sound isolating headphones on. The key takeaway of this sketch was the level of tracking fidelity needed to make the original goal of this project -- mid air button clicks -- work well.

What seemed simple at first glance required that I interface with the Oculus Constellation tracking system in the controllers, deconstruct and parameterize a series of waveforms, map those parametric waveforms to real world data, and then build a complex multi-sensory and multi-modal experience. Oculus and I decided to park this concept as a result, deciding to focus on the very compelling button augmentation and joystick augmentation concepts.

FINAL CONCEPT FROM THE SKETCHES (MICKEY)

Beyond all the haptic sketching, I also wanted to present a vision of what would be. If the effort was put in to optimize all the sketches, what sort of controller would come out? Mickey is the result of this question, the culmination of five weeks of research, prototyping, and validation. This design has three key features that incorporated the conceptual work from the haptic sketching period.

WIDEBAND HAPTIC ACTUATORS & LOFELT DSP

Rather than a single LRA, the Mickey relied on two wideband actuators, each a LoFelt L5 VCA. These actuators were much more powerful than the LRA, and were mounted orthogonally to one another, ensuring even propagation of the vibrations through the controller and into the hand for even the oddest competitive Beat Saber grip -- improving performance.

The LoFelt L5 is a voice coil actuator (VCA) with a resonant frequency at 60hz. The actuator works from 5 Hz to well beyond 1kHz, covering the haptic spectrum and a chunk of the auditory spectrum. This more comprehensive range of frequencies creates a more realistic and immersive touch experience. The lower resonant frequency of the L5 actuator allows haptic feedback that is in line with the perceptible tactile range of the skin. For human perception, low frequencies are vital to creating a pleasant experience. High frequencies above 250Hz can create numbness and tingling over a more extended period.

When comparing the L5 to the more generic Oculus Touch 2DLRA or Oculus Question Touch LRA, the experience of the Mickey is more realistic and immersive. The older designs much simpler "buzzes" fail to create compelling, immersive experiences because they do not match the high-quality audio and visuals featured in the IVE.

The integration of the LoFelt WAVE Digital Signal Processing (DSP), which detects the shape and frequencies of a game's soundscapes and converts them into high-definition (HD) haptic signals changes the haptic touchscape of the game. Rather than predetermined and limited buzzes and rumbles, the DSP, when combined with wideband Lofelt L5 haptic actuators, translates the signals into high-fidelity vibrations, bringing the IVE to life.

TWO LOCALIZED HAPTIC ZONES

A second key feature was the splitting of the controller into two distinct haptic zones. This separation would enable new effects from the motions and gaming category, adding new immersive experiences to the Oculus ecosystem.

To create a set of zones, I first had to figure out how to package the actuators into the controller and identify available space inside the controller. From online teardowns [13], and the teardowns I did in Milestone A of the project (see appendix A), I knew that two large voids existed in the Oculus Quest Touch controller. The first was the battery compartment, and the second began behind the grip trigger and extended to the main PCB -- just tall enough for an L5.

By cutting the controller just below the grip trigger, I could separate these two voids from one another. This also lined up nicely with the middle finger, allowing me to create targeted interactions to the "tripod grip." The proposed upper zone would react to the grip trigger, the main trigger, and take up the role of the traditional rumble motor. The lower area would respond to the joystick, X, Y/A, B buttons, and support any apparent haptic motion effects. The two zones would allow for new interactions -- improving immersion and realism. Additionally, when working together, they would ensure clean propagation into the hand, a performance issue of the current Oculus Quest Touch controller.

FORCE DEPENDENT BUTTON

To make room for both LoFelt actuators, the grip trigger had to be converted from a physical button to a force-dependent resistor. The basis for the concept goes back to an Oculus Patent [20] filed in 2019 by Jason Higgins and Benjamin Rogoza. That utilized a pressure sensing switch in place of the grip trigger's hall effect sensor. In the patent's claims, a series of methods are outlined to generate a hand animation in proportion the force inputs on the pressure-sensitive area.

In Mickey, I took the concept one step further. Rather than having a simple pressure sensor without augmentation, I thought it would be useful to couple the pressure sensor with a voice coil actuator. This pairing would rep-introduce some physical sensations into the pressure sensing actuator design, recovering the user interaction experience that was lost. By linking the two devices behaviors, I could re-create the feeling of physical travel. Another useful side-effect was that I could cue the users' behavior -- allowing the force-dependent resistor to feel like a hard surface or a deformable one. I could also add in the ability to make the grip trigger feel like various buttons -- a double click, single click, etc.

The most significant advantage of this change was in user comfort. I could mold a more comfortable shape that was permanent. The hope was that this would help clarify how to hold the controller for new users and help convey user intent. I could also provide a more consistent and controllable vibration to the finger, improving performance. Rather than rely on the button to transmit vibration--a somewhat risky proposition given the degrees of freedom intrinsic to a buttons workings.

BATTERY PACK & EXTENDED LENGTH

The Oculus Touch and Quest Touch controllers both use disposable AA batteries, typically a consumer pain point. Users expect their devices to work when they need them, and dead batteries can be extremely frustrating. Additionally, the disposable cells are harmful to the environment, if not properly recycled. Most consumers don't have access to or are unwilling to use designated battery recycling centers, so Alkaline batteries end up leaching chemicals while sitting in a landfill.

To address the weaknesses of disposable batteries, the Mickey concept uses a rechargeable battery pack in the grip. This mounting point at the very bottom of the grip has the side benefit of addressing comfort issues that came about when moving from external (outside-in) tracking to headset mounted (inside-out) tracking.

To be seen by the HMD, the LED tracking moved from a low slung orientation to a new raised position. The rearrangement drastically altered the center of gravity of the Oculus Quest Controller, leading to a front-heavy bias. During use, the Quest Touch produced more wrist strain and discomfort over extended interactions than the original Touch. A battery pack that slightly extended the length of the controller below the hand and redistributed mass towards this extreme would rebalance the controller, and increase user comfort.

Fig: Oculus Quest Touch superimposed over “Mickey” concept”

IMPLEMENTING INTERACTION-CENTRIC HARDWARE DESIGN

CHAPTER 08

IMPLEMENTING INTERACTION-CENTRIC HARDWARE DESIGN

At the halfway point in my thesis, I'd already done quite a bit. I'd spent hours with my nose deep into psychophysics and human perception research papers to help understand the body-centric side of haptics. I'd traveled to Berlin and Paris to meet with LoFelt and Actronika, trying to understand the limits of vibrotactile technology. I'd also formulated an approach that felt right to me -- the Interaction Centric approach -- hoping to fill the voids in haptic design thought that I'd experienced first-hand.

I'd spent hours deep in discussion with the entire supporting cast of this project -- my professors and advisors at the TU; other researchers across the globe; Oculus engineers, designs, and third party engineers and designers. At this point, it was time to put things into practice. The gauntlet had been laid down -- improve the Oculus Touch experience in VR. Deceptively simple, research revealed this problem to be a balancing act between four factors -- performance, comfort, immersion, and realism. I would also have to balance the human perception demands against the technical limits of vibration motor technology. The controller would have to be focused enough to create compelling experiences, but the actuators chosen would need to be flexible enough to create the full range of experiences the Oculus ecosystem demanded.

CHANGING PRIORITY: FROM BUTTON CLICK TO BUTTON AUGMENTATION

The midpoint also gave me a chance to reflect. Of the twelve proposed interactions, the primary focus had been to create a compelling button click in mid-air. This goal had proven somewhat elusive; the mid-air nature of the interaction wasn't well suited to the constraints of the Oculus controllers. However, In the haptic sketches, I'd explored lots of exciting and novel potential experiences, some of which were quite compelling. After talking with Oculus, we all agreed that mid-air buttons seemed a bit trickier than we'd expected. After trying all the haptic sketches together, we decided to pivot the project to prioritize the experiences from the Faline sketch, transitioning from button clicks to the augmentation of the buttons on the controller.

The experience provided by Faline was incredibly compelling. Even with simplified and abstract waveforms and only utilizing the tactile layer, there was an undeniable feeling across the team that this experience was really "something."

This change also reflected the interaction goals of Oculus. While mid-air buttons

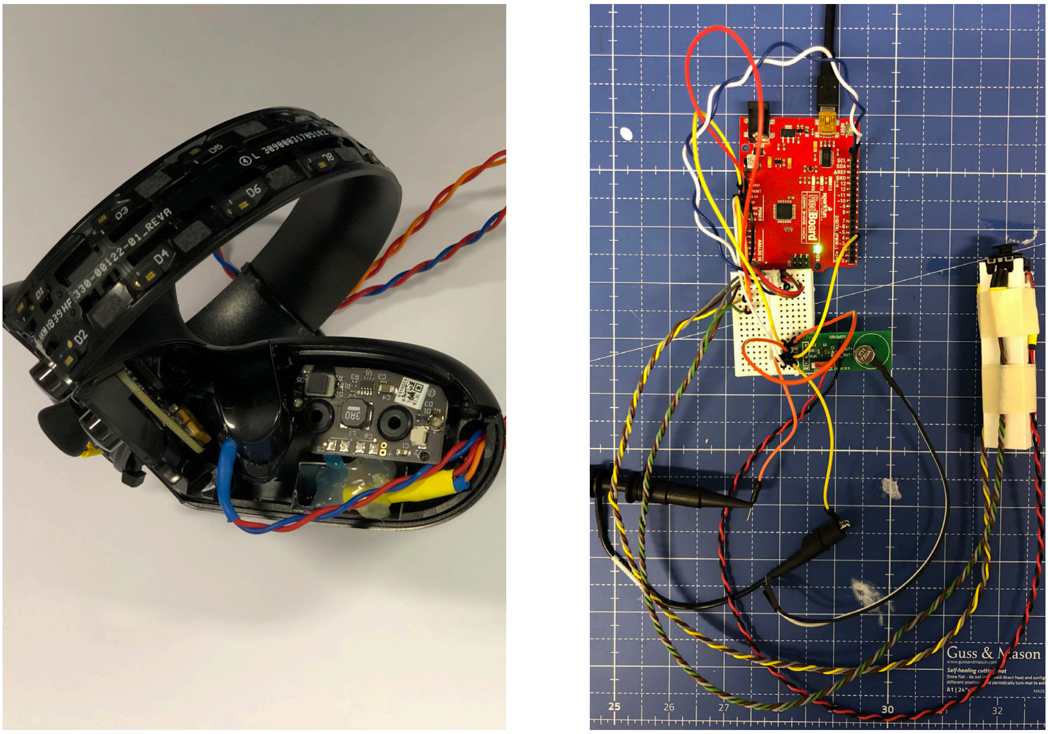

Fig: Sketch 01 and Sketch 08, chosen for further iteration

are a crucial target for improvement, the primary way of interacting in VR remains the physical buttons on the controller. Improving the experiences involving the physical buttons would create a more considerable overall improvement to the Oculus ecosystem.

We also decided to advance the Aladdin sketches to this final phase. I'd had a soft spot for these haptic sketches, and this was a completely new experience to the Oculus team. While they'd built various prototypes that created the experiences of most of the haptic sketches, the Aladdin sketches were wholly modern and novel to the team. My fondness for the sketches aside, we all wanted to see what would happen with a bit more effort and a few more iteration cycles. Additionally, Aladdin was possible with the existing hardware of the Oculus Touch and Quest touch controller, meaning if it worked, implementation would be relatively trivial.

This focused direction helped clarify the challenge. I now knew the scope of my final deliverable: one set of controllers, a left hand with Joystick Augmentation, and a Right Hand with Button Augmentation. Each would have unique challenges to solve, but if I were going to use my new Interaction-Centered approach, I would first I need to devise a demonstration/scenario that would allow me to explore the interactions.

SCOPING DOWN: SELECTING THE MOST COMPELLING SCENARIO

On the last day of my second Oculus visit, the Oculus team and I sat down to come up with several interaction scenarios. We wanted to find the experience or set of experiences that would make evident the benefits of these hardware changes to anyone who tried my final deliverable. While we weren't sure exactly what the hardware changes would be, we knew that any cost increase to the controller would be a hard sell to Oculus leadership. Their mission -- to bring VR to the masses --hinged on an aggressive pricing strategy.

The retail price of the Oculus Quest and Rift-S is $399. Looking at the teardowns I'd seen, I made a BOM cost estimate of around $150-$200 for either headset. Factoring in logistics, packaging, assembly cost, I figured the best-case scenario for the Oculus headset was a $100 in profit per set. This meant that even a $1 increase in controller BOM would be a noticeable percentage of the profit. Whatever demo we came up with would have to be compelling enough to justify that loss to have any chance of making it onto the product roadmap at Oculus. With the stakes set, we began brainstorming.

SCENARIO NO.1 SHOOTING GALLERY

The first idea was a fairly generic or common shooting gallery scenario. We'd put the user into a scene that let them shoot two guns, one semi-automatic, one fully automatic at various targets. The semi-automatic gun would allow the user experience pulling a trigger -- we wanted the ability to feel the mechanics of the gun, and ideally, use some of the trickery behind the Genie sketch to create a force effect. The fully automatic weapon would allow for some stronger, continuous vibrations. We also wanted to use apparent haptic motion to cue the user to experience recoil. The hope was that a slight timing difference between grip and trigger actuator would be enough to give the user an upward rise effect.

While we liked the idea, there were two significant issues. First, the shooting gallery experience was reasonably conventional. The Steam store and Oculus store -- the two leading marketplaces for VR games -- had dozens of similar apps. We wanted something that stood out, and a shooting gallery wasn't that unique. Second, the shooting gallery didn't highlight the Aladdin series of sketches. If I was going to incorporate joystick augmentation alongside button augmentation, the experience for testing should prioritize them equally.

SCENARIO NO.2 GRAB TABLE / SANDBOX

Our next idea came from an in-house only demo called the Sandbox. The Sandbox demo had a bit of reverence around the Oculus office, a template for the kind of successful outcome we wanted. While I wasn't able to try the sandbox demo, it was described to me as a sort of VR playground. The Sandbox contained a table with various virtual objects. The user could pick them up, manipulate, and otherwise interact with each object. What was also lovely was that such an experience didn't overly focus on the gaming market. It was something any IVE would require.

In our version of the Sandbox, we'd have a table with a few different balls. The user could pick up either a hard ball, a squishy/ deformable ball, and a breakable ball. We wanted to see if actuators in the buttons could make the user feel a different force effect -- either something rock hard, or soft and spongy. The breakable ball would go from hard to non-existent quickly, creating a result of something akin to breaking a glass in hand. Additionally, the user could move around the table and bounce the balls off of a wall, playing catch with themselves.

While this concept was an improvement on the first experience, we still didn't feel like the joystick augmentation was fully utilized. I also wanted a bit more purpose in the scenario -- with just a free play, I felt the scenario didn't take advantage of the satisfaction of accomplishing a task. I believed that a stronger emotional experience or journey would create greater engagement with the new controller design, and therefore better show the potential of the new design.

SCENARIO NO.3 SLINGSHOT

The third idea drew from the popular game, Angry Birds. Instead of doing a shooting gallery with guns, we thought maybe a shooting gallery with a giant slingshot would be fun. The user could steer the slingshot using the joystick, and then grab a large elastic band and pull back. The pulling would allow users to feel the tension inside the elastic through vibration. The further they drew, the stronger the vibration would be. They'd be manipulating objects, feeling object material properties, and by using the joystick to steer the "turret" of the slingshot, they'd be able to feel friction.

This experience once added in a more "fun" element than our prior engagements. Rather than some abstract series of interactions, we felt the "Angry Birds" style slingshot game might be relatable to existing and popular games -- allowing for a quick interpretation of the improvements. The intent behind this game mode was to create an experience that felt relatable and clearly showed the improvements that the new actuators could provide.

However, this scenario still didn't meet our needs. It didn't fully take advantage of the joysticks, which I really felt were the novel part of the experience. I wanted to show increased immersion, and this scenario didn't provide that. I believed that the scenario needed an avatar that moved through the digital world.

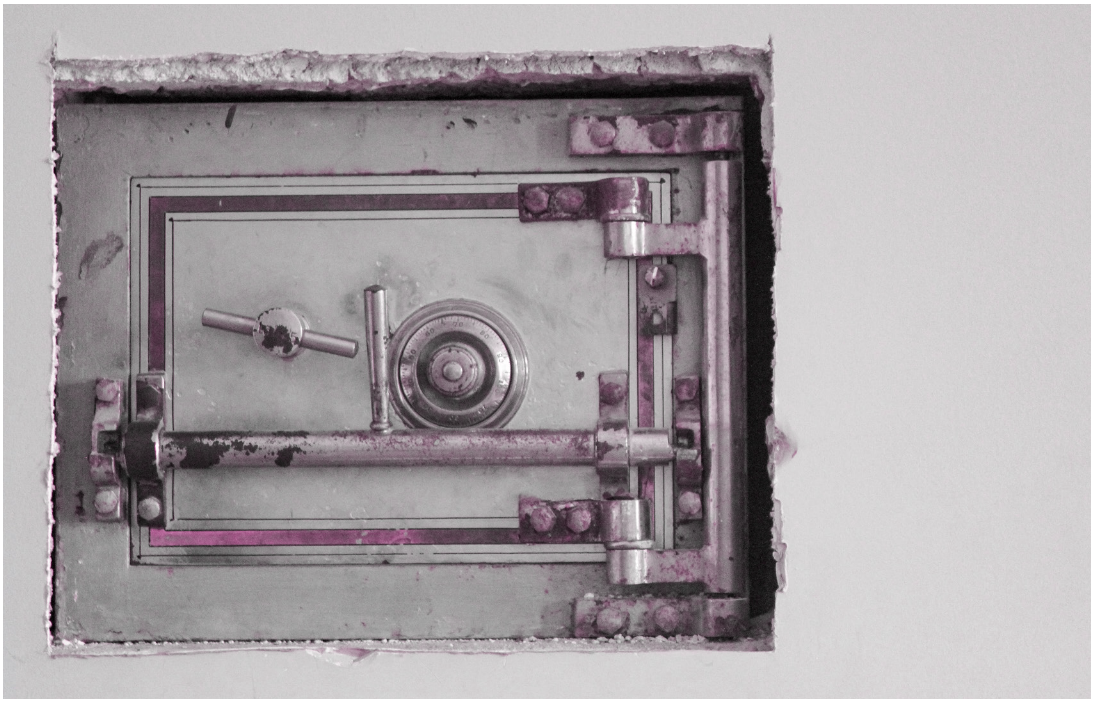

SCENARIO NO.4 SAFECRACKER

For our fourth idea, I wanted to focus on the task. I tried to think of an experience that would be difficult, entertaining, and rely on touch, I thought back to Skyrim, a role-playing game in which you picked locks through the joysticks. By applying the right amount of force of the left stick, and correctly positioning your lock pick with the right joystick, chests would open.

If we added in additional challenges, we could demonstrate the full feature of both button augmentation and joystick augmentation. Rather than do lockpicking, we also wanted to do a more stereotypical safe crack, with users having to decode the combination of a dial through touch. The hope was that very slight changes in vibration while turning the dial could be felt through the new augmented buttons. There would also be a lockpick similar to Skyrim's that would take advantage of the joystick augmentation.

This concept felt to us like the best idea so far, but there were still a few missing elements. One of the most significant advantages of joystick augmentation to me was in movement. If we were going to pick a static safe, we'd be missing out on what was to me, the most compelling part of the joystick augmentation experience. Also, safecracking could be incredibly frustrating if not done well. There was a chance the experience would be too difficult, going from a challenge to an unplayable frustration.

SCENARIO NO.5 ZEN GARDEN

Our final idea came about as a lighthearted joke. In looking at all our prior approaches, we wanted something that was not frustrating, and as far away from the gaming experience as possible. To us, that was a Zen Garden. We would put the user into a Zen Garden full of sand and surrounded by an asphalt ring. The user's IVE avatar could walk around, with the joystick controlling movement. The user could stroll, or if they moved past 50% of the displacement of the joystick, they'd begin to run. The joystick augmentation would create a force perception akin to the increased effort of moving about in the sand. We also wanted to play footfalls through the actuator. We'd also give the user a rake to rake the sand into various shapes and textures. The button augmentation would use a bit of apparent haptic motion to provide the rake with the "tugging" against the hand that the frictional force from the sand would create. We'd also have various boulders and objects that could be picked up or pushed against.

This experience seemed to hit most of the goals we had. It was task-oriented without being frustrating. It showcased both joystick and button augmentations. It was fun, interesting, and compelling. Most important, it was the first idea that any of us would be willing to pay for -- if it had been a game, we'd buy it. The Zen Garden felt right, and with my flight back from Seattle coming up, we'd run out of time on the brainstorm session. It was time to break down the Zen Garden and deconstruct the experience.

BREAKING DOWN THE ZEN GARDEN

THE EXPERIENCE

The experience level is all about user context and user experience objectives. The user experience goal of the Zen Garden experience is to produce a state of relaxation in the user. I want users to have the same tranquility and relaxation they would experience from a traditional Japanese Zen Garden. If I could have users enter a meditative state, and feel contemplative, relaxed, or tranquil, the experience would be a success.

I could rely on the cliche of the Zen Garden to help steer the user to the correct mindset. The act of raking the sand into a pattern recalling waves or rippling water, known as samon (砂紋) or hōkime (箒目), is understood even among non-Japanese as a meditative or contemplative act. I could reinforce this with other traditional proto-Japanese "Zen" sounds and visuals.

THE INTERACTION

The Zen Garden scenario would have four main interactions. The first would be the act of walking around the Zen Garden. This interaction would have an Asphalt mode and a Sand mode. This interaction would demonstrate the potential of the joystick augmentation.

The second interaction would be the user picking up and setting down the rake. This interaction would require the user to pick

Fig: Zen Garden VR experience Screenshot

up the rake from its resting place on the ground using the digital hand model and the grip trigger and trigger.

The third interaction would be the user raking the sand. While holding down both triggers, the user would drag the rake through sand, creating a visual displacement and pattern.

The fourth interaction would be the user hitting a rock with the rake. If the user hit the stone too hard, the rake would have to break.

THE EVENTS

For the first interaction (footsteps) the events would be fairly straightforward, a footfall each time the user took a step.

The second interaction would have three events. The first would be the user pickup event, the hand going from empty to full -- with a weight. Here we would have to create hand presence. The second would be the user holding onto the rake. We would need to create some sensation of weight in the hand, and the inertia of the rake in the hand. The third would be the release or drop event, with the weight going away.

The interaction also had three events. There was a ramping into the sand, dragging through the sand, and a

THE EVENTS (CONT'D)

release from the sand. All would have the same sort of effects, but with different ramp-up/ramp-down phases.

The fourth interaction would draw some events from the second interaction. The user would feel the rake in their hand and have the ability to swing. The contact with the rock would be the only new part of this interaction, with the impact effects of the collision.

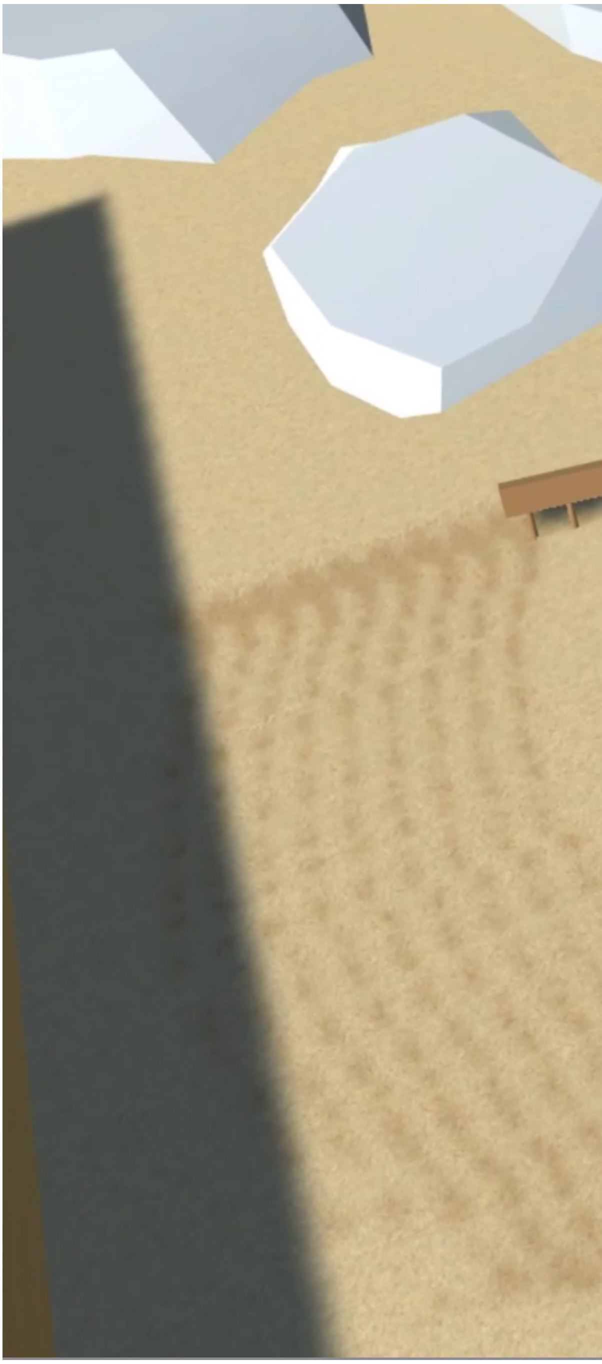

THE EFFECTS

Once again the footsteps would be fairly straightforward. We would need to create an effect of impact, and create an effect of motion. These would have to be parametric, changing for each of the surfaces the user walked over to demonstrate different deformations and energy transfers. The visual layer would handle most of the movement effects, with the user's point of view (POV) moving through the environment. We could further augment the impacts with auditory layer cues in conjunction with the tactile layer cues.

For the rake pickup, we would have to create an effect for the weight in the hand, the effect of moving the rake / stick slip in the hand, and the release of the rake. Once again we could rely on the visual layer to do most of the heavy lifting in regards to movement. Stick slip interactions would be the most difficult, but could be done through creative use of the apparent haptic motion effect.

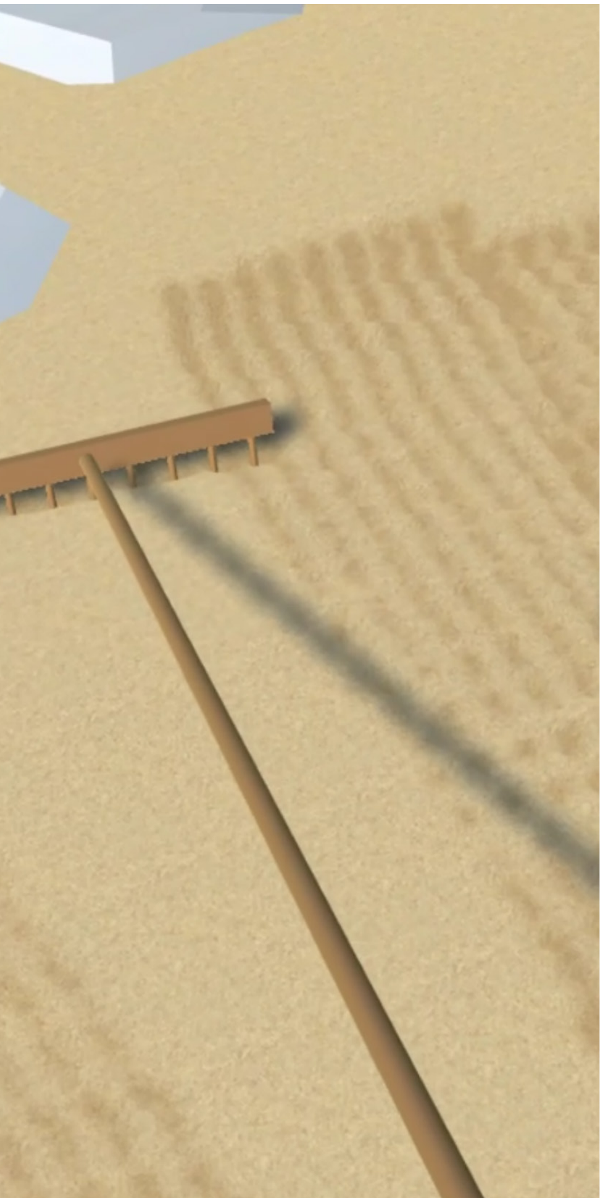

Fig: Raking sand in Zen Garden

The effect of the rake interacting with the sand would be the most complex task. We would have audio layer cuing about rake contact and release, but the ability to transmit second order tool effects is notoriously tricky. We would need to convey an effect of stick-slip and friction as the rake pulled through the sand, and the timing of this effect would be key to the realism and immersion of the scenario.

THE PERCEPTIONS

The perception of the footstep was pretty straightforward, with footfalls being a wonderful example of a simple haptic interaction. The rest of the interaction would require multi-sensory integration.

Using the control/display ratio we could help simulate the inertia and friction, creating a perception of the rake being either pulled out of the hand or lagging behind the rest of the movement. The C/D ratio refer to the artificial delay between the physical world and the digital world. By presenting visual layer cues slightly behind the real world counterparts, we perceive either friction or inertia. This visual layer cue can be augmented with tactile layer perceptual cues -- namely in the use of haptic motion effects and high frequency noise to produce frictional feedback.

The perception of impact can be presented visually and augmented with an impact waveform.

THE EVENTS (CONT'D)

THE WAVEFORMS

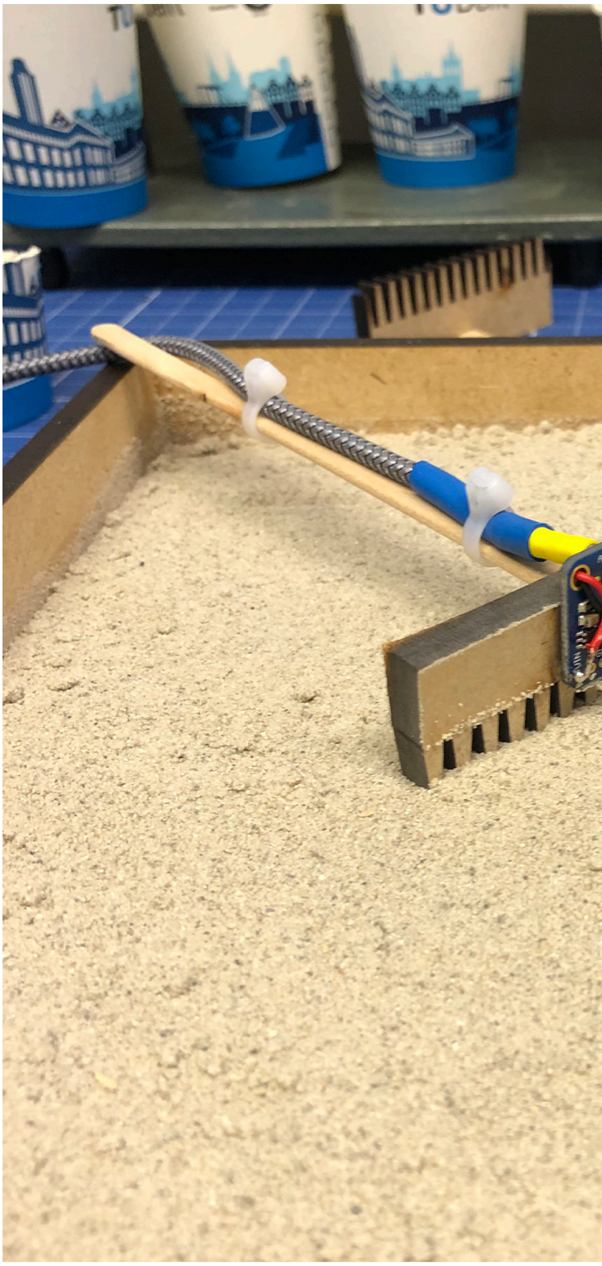

Since the actuators in the inTouch would all be driven by the LoFelt EVK, I needed a sound file as an input. The EVK would be run in haptic bypass mode, so there would be little interference from the DSP. The key was to get accurate waveform recordings, as the EVK would be replicating the exact output from the Unity Zen Garden environment.

For the footsteps, the sound-based approach was used. A microphone was strapped to my leg, and a series of waveforms recorded as I walked on asphalt and on sand. This loop was used to feed data to both the auditory and tactile layers of the Zen Garden experience.

The waveforms for the Zen Garden itself were created using the ADXL 345 accelerometer. Just like the button box recordings, the ADXL recordings would be used to produce sound files from the 3D acceleration data. The ADXL 345 was mounted onto a rake which was then used to rake sand in a desktop Zen garden.

The collisions were also done this way, with the ADXL 345 mounted onto a long pole, which was swung into a hard surface at three different speeds.

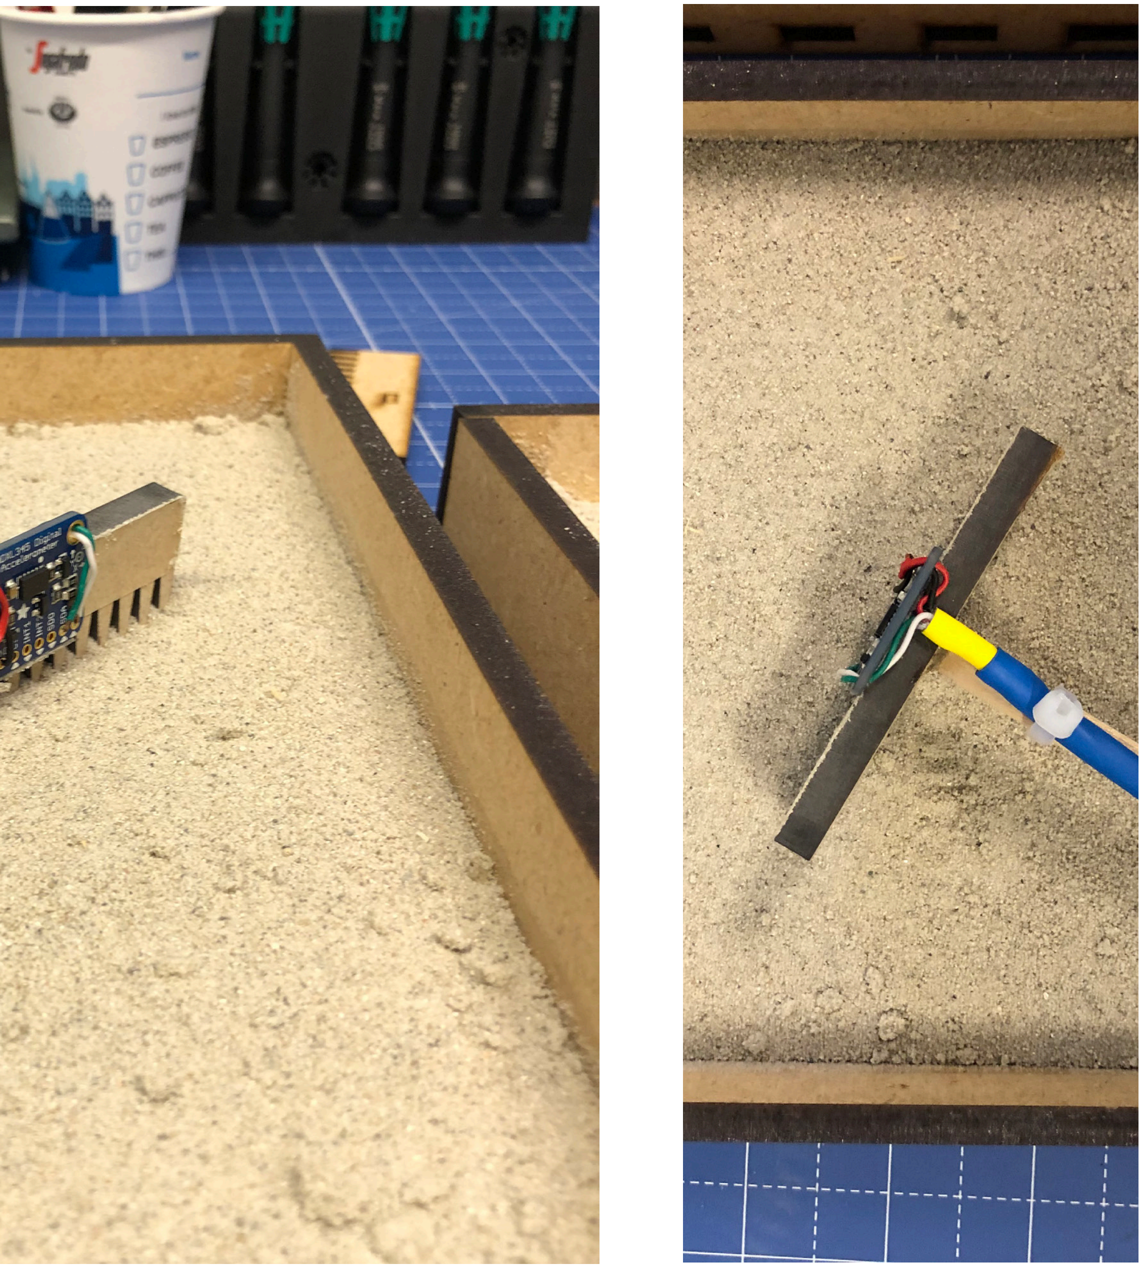

Fig: ADXL345 attached to rake to record haptic feedback and generate a waveform for the Zen Garden Experience.

SELECTING HARDWARE TO CREATE THE INTERACTIONS

With the interaction well understood, I could begin to explore what sorts of actuators would be needed to achieve the desired effect. I began with one final haptic sketch.

SKETCH 11 - BUTTON AUGMENTATION II (CINDERELLA)

The proverbial belle of the ball, Cinderella was the final haptic sketch 'I'd build under the scope of this thesis. Cinderella was constructed to test actuator orientation and style inside the grip trigger and main trigger of the Oculus Quest Touch controller. Building from the original CAD provided by Oculus, the buttons were modified with a series of Jin Long G0832013D coin cell LRAs.

The grip trigger, which sits on a lever arm inside the controller, was nicely vibration isolated from the rest of the controller body due to the additional degree of freedom of the lever arm. By removing some supporting structure inside the grip trigger, it was easy to incorporate a mounting fixture for the LRA. Some small routing for wire strain relief was incorporated, before sending the CAD off for manufacturing. The LRA fired perpendicular to the radius of the lever arm, rather than along the axis. This orientation isolated any vibrotactile effects from the rest of the controller, as the lever arm and its return spring worked as a natural mass-damper system.

The main trigger was a bit more challenging to design. The most effortless orientation for the integration was axial, unlike the grip trigger. When trialed in this position, any vibrotactile effects in the trigger were compelling, but they were not isolated from the main body of the controller. I wanted a clean, crisp haptic signal, and didn't want additional noise from the resonances within the controller body. A second iteration rotated the LRA to be perpendicular to the swing of the grip trigger, producing the same clean isolation as we saw in the grip trigger.

With orientation determined, it was time to select the final actuators.

Fig: Physical prototype to determine button changes

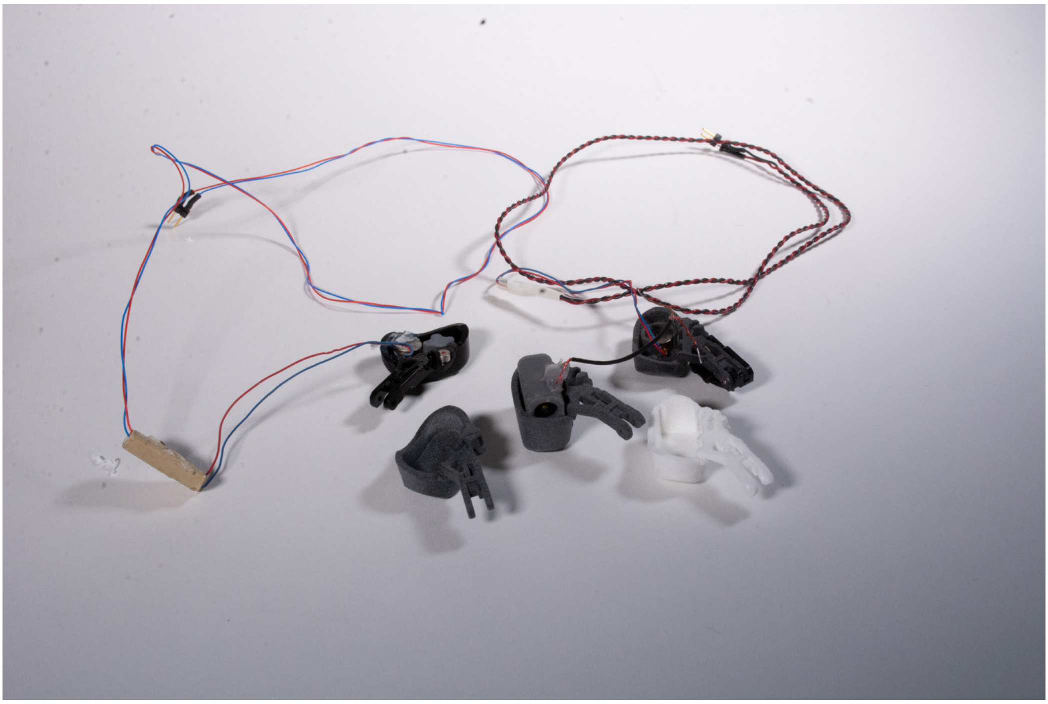

Fig: Button Interations for the Grip Trigger feedback.

Fig: Button Interations for the Trigger feedback.

JOYSTICK ACTUATOR SELECTION

From the breakdown, it became clear that any actuator or actuator combination would have to operate in two frequency bands. The first, the "Haptic High Frequency" (400Hz-1kHz) range would be vibrations for friction emulation. This high-frequency noise was demonstrated first in the Genie Haptic Sketch. I also needed an actuator that could produce the footstep vibrations which occurred in the "Haptic Core" Frequency (100Hz-400Hz). There were three options to create these vibrations -- a pair of LRAs, one of which would need to be customized for the high-frequency noise, piezo actuators, or a voice coil actuator.

The LRA pair would be cheap. LRA technology was well understood, and the driving would be reasonably straightforward. Although using two LRAs would require additional driving circuitry, at the production volumes of Oculus, this would be a negligible expense. What concerned me was the complexity of such a system necessitated. The logic behind controlling the different LRAs would be complicated, and implementation for the software developers would be a bear. Furthermore, all this complexity would still be limited to two narrow bandwidth ranges.

The piezo implementation would be the easiest to package. Piezo haptic elements are incredibly thin, meaning the actuator could be located inside the joystick itself. This super close co-location would minimize the mass driven by the actuator, thus reducing the actuator size and impulse requirements. Additionally, the piezo element could be used as an input device, removing the need for a joystick that contained a click button. Unfortunately, piezo technology isn't quite ready for mass/consumer market. The individual actuators would be too pricy for the Oculus controller. Long term, this would be the ideal option.

Fig: Base Actronika Haptuator

In contrast to the LRA, the VCA would be a single actuator that could hit both of my required ranges and still have some headroom to reach lower frequencies -- useful for everyday VR experiences such as explosions and shooting games. VCA technology is just at the point of maturity to be viable. Packaging the actuator inside the controller would be a challenge.

With the style of actuator selected. I had to choose the actuator itself. I had three main choices: the LoFelt L5 actuator,

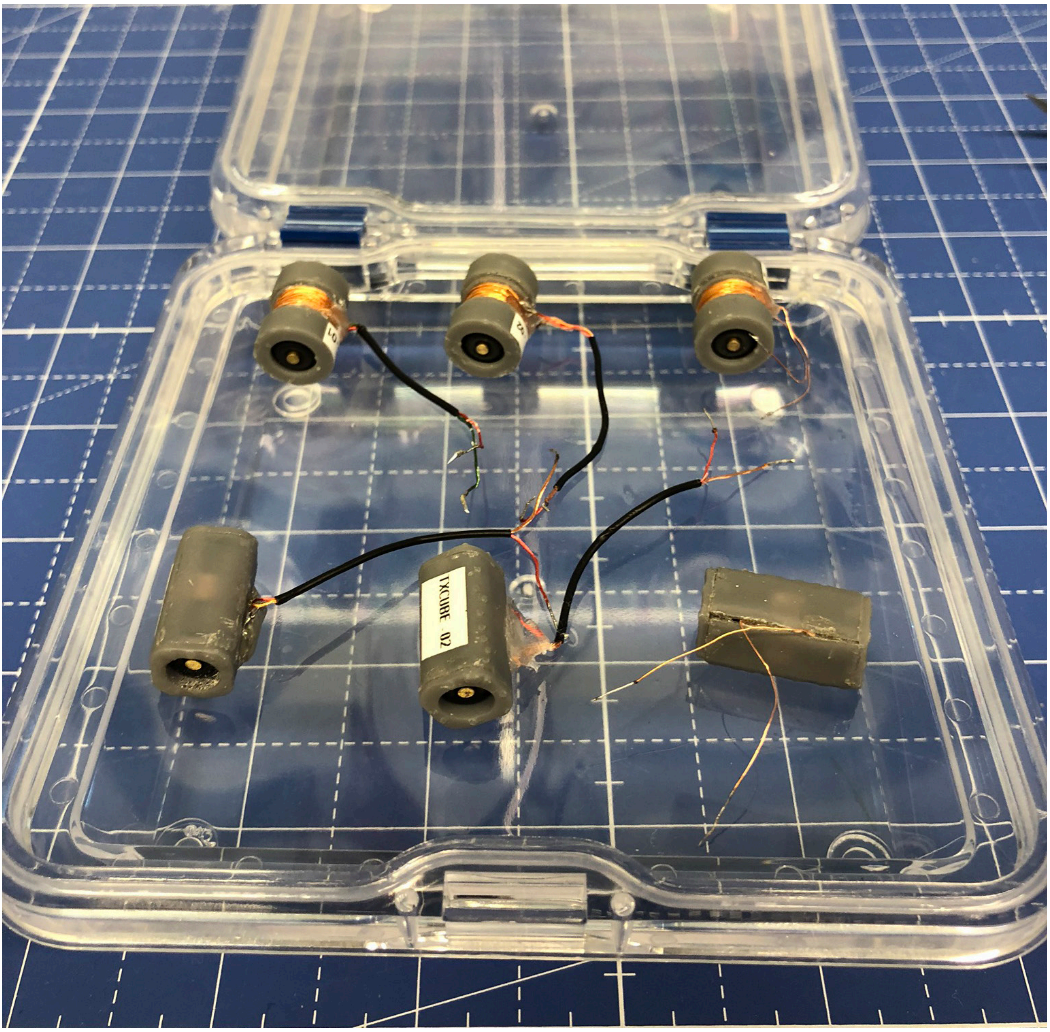

Fig: Custom Actronika Haptuators. Front for Grip Trigger, rear for Trigger.

the Actronika Haptuator, and a custom-designed NIDAC actuator Oculus had provided me. All had similar impulses -- the NIDAC being slightly weaker and the Haptuator being a bit stronger. The decision would come down to the packaging. The LoFelt was a short rectangle, the NIDAC a cube, and the Actronika a long cylinder.

The Actronika would fit nicely in the battery compartment of my prototype, and short-term would be the best option. In a longer term design, any of the three would work, and the controller would need to be designed to accommodate the actuator chosen.

Due to the nature of the waveform generation -- it was coming from the Unity VR scene of the ZenGarden, audio to haptics would be required. While it was possible to custom design a DSP and driver board, due to the time limitations, I selected the LoFelt EVK as the driver.

BUTTON ACTUATOR SELECTION

As with the Joystick augmentation, any actuator or actuator combination used for the augmentation of either the trigger or grip would have to operate across multiple frequency bands. Due to the complex nature of the potential interactions that would be felt through the grip trigger and trigger, a wide-band actuator was a must. Much like for the Joystick Augmentation, Button augmentation would require the use of a VCA.

The challenge for packaging of a VCA would be much higher for the button augmentation effects. The joystick augmentation could be one with a single actuator, but button augmentation required two unique actuators. These would each need to fit inside the respective buttons and also not cause haptic interference/ noise with the other.

Even with all the contacts, I'd made across the industry through my thesis or my work at SenseGlove, I knew of no actuator small enough to meet these requirements. While LoFelt had a small actuator called an L3, it was very much a prototype, and not anywhere near market-ready. I pivoted from searching for an actuator that would work, to exploring the fundamental concepts behind the actuator. The goal was to find a design that could be easily miniaturized to fit inside each button.

All VCAs work much the same, there's a coil, a spring, and a permanent magnet. The design limit for most actuators is in the spring. Each company used a unique approach to the rebounding spring in the design, and whichever spring mechanism would scale the best would likely be the actuator that would scale the best.

I again begin by looking at the LoFelt concept. The design of the LoFelt relies on a brass sheet to work as the return spring. The moving mass of the VCA is mounted on one end of the brass "dogbone" shape, with the other end attached to the casing of the actuator. Much like bending a plastic ruler off a desk and letting it "snap" back, the dogbone pushed back against any deflection. The spring force was then a function of the material, the length of the "dogbone" strip, the cross-sectional area of the strip, and the deflection of the strip. The use of two springs -- one on each side of the casing -- balanced out the effects of the spring force on the movement of the actuator, allowing it to effectively free float.

This design wouldn't scale easily. The issue I had with the LoFelt was the length of the axis on which the strip was mounted. While super easy to manufacture such a spring, finding a shorter throw equivalent would require a lot of trial and error with the material selection and geometry. The L5 design wasn't what I wanted.

I then turned to the Actronika Haptuator. Their flat series looked promising but required laser welding and sheet bending to produce the casing. These tooling needs meant further involving industry partners, and lengthening the timeframe. While the flat series wasn't going to work, their standard profile seemed promising. The main issue with the haptuator design was length down the axis of movement, something that was easy to scale. The spring mechanism of the haptuator was a toroidal thrust bushing made of an elastic material. This design is very space-efficient, so provided that I could produce the required forces from a smaller number of windings the haptuator would work.

I reached out to my Actronika with my requirements and asked it would be possible to turn their haptuator into a "halftuator" -- same form factor with a shorter length. Together, we came up with two designs that would just barely fit in the button profiles on the controller. I would need to redesign the buttons to accommodate the actuators, but I wouldn't need to redesign the controller itself, an essential requirement.

Actronika delivered two haptic actuators rapidly prototyped through the Formlabs 2 stereolithography printer. The coils were hand-wound, causing some noise in the output signal, but the actuator was able to deliver precise, compelling vibrations from 8Hz-15Hz. The maximum output wasn't specific, but to me felt in the range of 3.0-4.0G on a 100g mass, more than enough for any effect I wanted to produce.

Due to the nature of the waveform generation -- it was coming from the Unity VR scene of the ZenGarden, audio to haptics would be required. While it was possible to custom design a DSP and driver board, due to the time limitations, I again selected the LoFelt EVK as the driver.

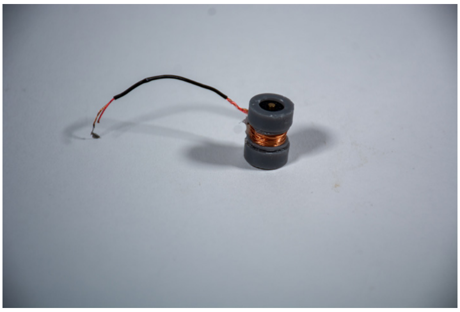
Fig: Haptuator used in grip trigger


Fig: Grip Trigger at various depths and actuators

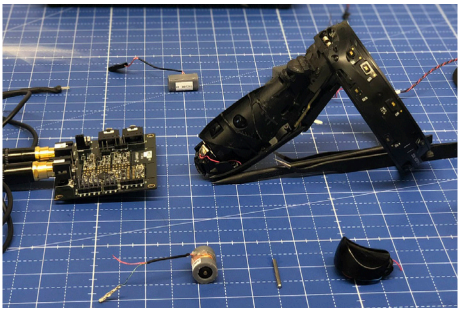
Fig: LoFelt EVK (left) Haptuators (center), modified Quest Touch (right)


Fig: Trigger prototypes with various actuator orientations and actuators.

THE FINAL INTEGRATED PRODUCT

CHAPTER 09

THE OCULUS QUEST CO
(BASELINE DESIGN)

The baseline for the Oculus inTouch prototype is the Oculus Quest controller. Any improvement made would have to fit within the plastics and confines of the preexisting Oculus Quest Touch design.

The key features of the Quest Touch controller includes: an analog thumbstick similar to those found on other modern game controllers, two face buttons that can also be pressed with a thumb, a trigger designed for the index finger, and a second trigger that is activated by squeezing the rest of the fingers against the controller grip. Each of the triggers and buttons contained a number of capacitive sensors that are capable of telling where the player's fingers are located, and were used to generate the digital hand model.

The design for these controls worked well, with each button ergonomically located. From a haptics perspective, there was no need to change any of these buttons. The design and placement of these controls was well optimized, and balanced the four design factors. These buttons then became a locked part of the design. As much as possible I would avoid changing or modifying any control input location.

Another locked part of the design was the constellation of IR LEDs to track the position of each controller. The ability to have accurate hand tracking is key to the design. If the real world movements don't correspond to the digital world it wouldn't be possible to create any sort of immersion.

NTROLLER

Fig: Oculus Quest Touch controllers.

OCULUS INTOUCH: THE FEATURES

Fig: Oculus inTouch controller exploded view, front quarter.

The Oculus inTouch controller contains five key design features:

Redesigned Joystick: The inTouch controller allows users to feel closer to their digital avatar than ever before. The new joystick will enable users to sense every footstep and experience every surface. High-frequency waveforms make users aware of the effort required to maneuver in digital space, tricking the brain into feeling a resistive force where there is none. The new joystick helps improve two design factors -- realism and immersion.

Redesigned Trigger: The inTouch controller makes first-person shooters more compelling than ever. Feel the movement of the trigger, and the thump of the recoil in your hand. New button colors, materials, and finishes, (CMF) help guide the hand to the proper positioning. A unique matte texture and semi-soft coating help keep the finger in place through even the most energetic vibration. The new button improves immersion and performance.

Redesigned Grip Trigger: The inTouch controller makes digital objects more believable. When grabbing objects, the grip fires back. New button colors, materials, and finishes, (CMF) help guide the hand to the proper positioning. The matte texture and semi-soft coating shared

Fig: Oculus inTouch controller, rear quarter view.

with the trigger help keep the finger in place and allow users to feel digital objects in hand better than before. The new button improves immersion, comfort, and performance.

Voice Coil Haptic Actuators: The inTouch controller takes advantage of cutting edge haptic motors. Custom designed by Actronika SAS, the triggers, joystick, and button all respond to every click, bump, and action with fidelity. The actuators respond from 8Hz well beyond 15kHz, covering haptic and audio spectrum ranges. With the ability to produce waveforms that fluctuate in frequency and amplitude, these actuators generate a more compelling interaction than ever before.

New Centre of Gravity: With the introduction of inside-out tracking, the Oculus Quest touch had to shift the tracking ring from around the hand to above the controller. This raised position of the tracking ring causes a bend in the wrist, which over time can cause a comfort issue for the user. The new VCAs and weighting at the base of the controller help bring that orientation back into alignment, improving overall user comfort

OCULUS INTOUCH NEXT STEPS

The goal of the Oculus inTouch controller was to produce a significant improvement in four design factors -- comfort, performance, realism, and immersion. Based on these criteria, the inTouch prototype is a partial success, and will require additional iterations. The inTouch prototype design is a stop-gap design that makes compromised to fit into the original Oculus Quest Touch controller design. These hard requirements limit the amount of improvement that can be made in the four design factors.

The inTouch design prototype little improvement in comfort. The virtual design of the inTouch controller and the prototype diverge. The inTouch prototype controller does not implement any of the postulated design improvements for comfort. Although the redesigned weight balance can be modeled due to the weights added into the controller, a more comprehensive set of changes is needed to actualize any true benefit in this design factor. The inTouch controller, much like the Quest Touch controller could use additional cues for use, donning, etc. The controllers, while not complex, often wind up in the wrong hands or orientation of first time users. Stronger color pathways and contrasts could be used to clearly illustrate where hands/fingers go. Additional exploration of materials and textures would greatly improve the comfort -- increased airflow to areas of the hand in contact with the surface of the controller would greatly reduce sweat.

The inTouch prototype does provide improvements to the performance of the controller. The ability to produce precise, repeatable and accurate interactions is improved through the new VCA designs. The ability to provide closed-loop control over an actuator that is able to affect a wide range of frequencies helps ensure that any waveform produced behaves the same each time it's presented. The more central location of the VCAs compared to the Oculus Quest LRA reduces distance between hand and actuator, reducing resonances and secondary vibrations within the assembly of the controller. The location within the button makes haptic feedback more crisp and compelling than the extreme top of the controller. What is needed now is refinement in the design. Additional iterations in button shape, orientation of vibration motor, and button hinge design are needed to help better isolate the VCAs from one another. Additional work on PCB design and DSP would help streamline the waveforms being presented to the controller.

Above all, to improve performance, a dedicated set of standards for VR haptics development are required. The linkage from the digital world to the real world is poor in the haptics realm as the data flow is not well understood. The ability to link digital world behaviors to specific actuators is non-existent, as most conventional controller designs only use a single actuator. The ability to map vibrations from digital world to single wideband haptic actuators in either hand isn't developed, much less the more complex control schema required to drive something such as the inTouch concept. Even a simpler version, such as the Cinderella concept which only uses LRAs, requires massive improvements to Unity, Unreal and other VR development environments.

The focus of the inTouch concept was to improve realism and immersion. These two interrelated design factors took precedence in my decisions, causing sacrifices to comfort and performance. On this the inTouch concept does provide an improvement. Real world interactions are not limited to narrow bands in the frequency domain, instead playing across wide bands of the spectrum. The use of VCAs that can output signals that more closely match the real world provide a massive boost to realism. The ability to tune time, amplitude and frequency adds another dimension to each haptic feedback. Realism is also improved through the greater spatiotemporal control over the vibrotactile feedback. By having two distinct and well isolated locations for vibration, we can more realistically render interactions closer to the point of digital interaction. Buttons now click at our fingertips rather than in the palm.

Realism is also improved by the new interactions made possible through the inTouch design. The inTouch controller produces a more accurate portrayal of the twelve interactions outlined in section "The Scope: Prioritizing and Interaction". Rather than approximations of each interaction, and abstractions presented through an LRA, the VCAs of the inTouch design make for waveforms that more accurately represent the behaviors of each of these interactions.

Immersion is also improved in the inTouch design. The improvements to immersion start at the approach to the design. The use of the interaction-centric approach behind the design places immersion at the forefront of the inTouch controller prototype. The understanding of the objective behind an interaction ensures that the actuator selection is done with specific sensations in mind. The VCA's size, power, and spectrum band allow for flexibility -- allowing the inTouch to provide a greater range of haptic cues. The advantage of haptic cues is that information provided to the user takes into account the relevance of information embedded in incoming signals/stimuli/information and the information reduction is achieved by user goals and predictions for the interaction experience.

The success or failure of the inTouch's ability to improve the experience in realism and immersion comes down to the ecosystem. Without developers and environments that carefully select what haptic information is presented, the inTouch concept loses it's effectiveness. A series of best practices and easily implementable approaches are critical. The use of LoFelt's sound-to-haptics approach drastically simplifies and streamlines the development of haptic cues, but the filtration algorithms used are not yet advanced enough for the mass market. Even the best hardware implementation will fall flat without the support of good software.

REFLECTIONS AND FINAL THOUGHTS

CHAPTER 10

REFLECTION

This project represents the culmination of the past two years of my life. While not a direct output of everything I've learned at the TU Delft, it is the byproduct of a new way of problem-solving and a new way of thinking. I came to the Netherlands wanting to work with and study how we interact with the world around us. I wanted to help create something that makes a difference in the lives of people and helped shape the way that we interact with ourselves, each other, and the world around us.

I came searching for more tools in my mental toolbox. I had in my bachelors studied how to build and create, I'd learned the engineers' approach to problem-solving -- the use of logic and problem deconstruction to develop solutions. I graduated and moved on to consulting, spending five years learning to explore the justifications and needs behind my creations -- the constant questioning of why rather than how.

TU Delft has given me the ability to explore not just the why's behind the product, but the why's behind the user. This master's thesis and the interaction-centric approach that acts as the cornerstone to work is a reflection of that new focus. The outcome -- more than just a physical product -- is an approach to designing immersive and compelling experiences that put the user's involvement at the forefront.

The evolution of digital spaces from 2D to 3D will radically change interaction paradigms. Rather than rely on input devices -- mice, keyboards, touchscreens -- devices that require users to interpret their desires into a format the computer can understand, we're beginning to see the first interaction devices. Interaction devices rely on intuitive, natural interactions that put the onus on the computer for understanding the user's desire. Touch will play a central role in this change. Haptics is an old field, that is undergoing a revolution. We no longer rely on haptic feedback as a notification or a novelty; it's now a core part of creating an engaging and immersive experience.

The next generation of these input devices will drive this home. The goal of the inTouch controller design is to help lay the groundwork for these still-theoretical interaction devices. The ability to create straightforward, understandable, and compelling interactions is going to be critical. I hope that future researchers, academics, designers, and engineers can find inspiration in this work, and will continue to grow and evolve the Interaction-Centric approach to haptic design.

APPENDICES

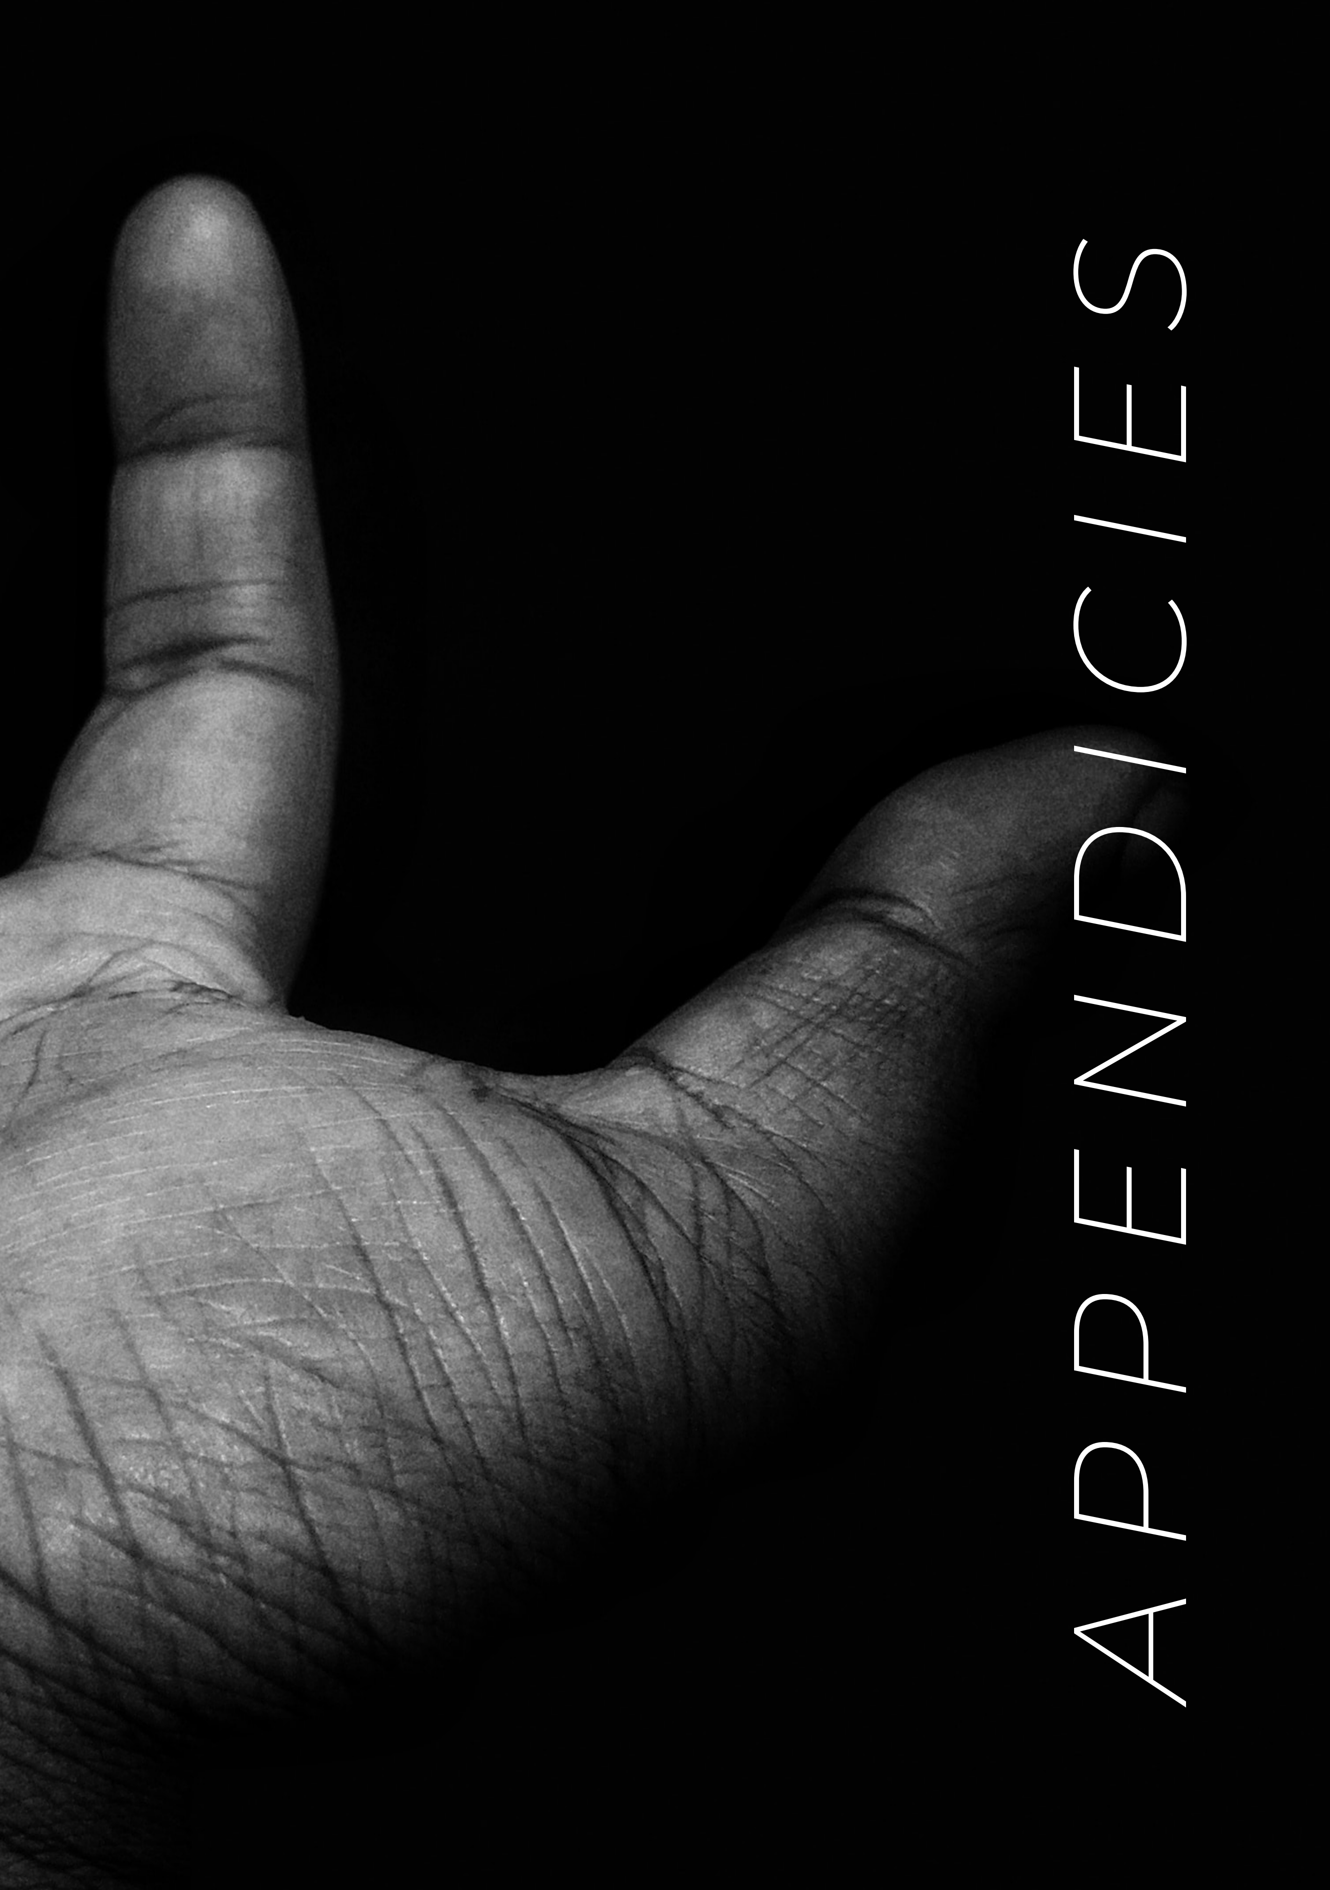

WORKS CITED

[1] D. Shor et al., "Designing Haptics: Improving a Virtual Reality Glove with Respect to Realism, Performance, and Comfort," Int. J. Autom. Technol., vol. 13, no. 4, pp. 453–463, Jul. 2019.

[2] D. Mallard, "Developing a Product Requirements Document for Hardware Startups," Medium, 25-Jul-2018. [Online]. Available: https://blog.bolt.io/one-of-the-most-powerful-tools-for-hardware-startups-writing-25cd261dbe0b. [Accessed: 18-Jul-2019].

[3] A. van Boeijen, J. Daalhuizen, R. van der Schoor, and J. Zijlstra, Delft Design Guide: Design Strategies and Methods. Bis B.V., Uitgeverij (BIS Publishers), 2014.

[4] E. H. Weber, E.H. Weber on the Tactile Senses. Psychology Press, 1996.

[5] M. Heidelberger, Nature from Within: Gustav Theodor Fechner and His Psychophysical Worldview. University of Pittsburgh Pre, 2004.

[6] M. Grunwald, Ed., Human Haptic Perception: Basics and Applications. Birkhäuser Basel, 2008.

[7] C. Starr, C. A. Evers, and L. Starr, Biology: Concepts and Applications. Thomson, Brooks/Cole, 2006.

[8] L. A. Jones and S. J. Lederman, Human Hand Function. Oxford University Press, USA, 2006.

[9] J. Dosher and B. Hannaford, "Human Interaction with Small Haptic Effects," Presence Teleoperators Virtual Environ., vol. 14, no. 3, pp. 329–344, Jun. 2005.

[10] R. S. Johansson, U. Landstro¨m, and R. Lundstro¨m, "Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements," Brain Res., vol. 244, no. 1, pp. 17–25, Jul. 1982.

[11] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," Nat. Rev. Neurosci., vol. 10, no. 5, pp. 345–359, May 2009.

[12] L. A. Jones, D. Held, and I. Hunter, "Surface waves and spatial localization in vibrotactile displays," in 2010 IEEE Haptics Symposium, 2010, pp. 91–94.

[13] M. Wiertlewski, J. Lozada, and V. Hayward, "The Spatial Spectrum of Tangential Skin Displacement Can Encode Tactual Texture," IEEE Trans. Robot., vol. 27, no. 3, pp. 461–472, Jun. 2011.

[14] L. A. Jones and S. J. Lederman, Human Hand Function. Oxford University Press, USA, 2006.

[15] V. Hayward, "Haptic shape cues, invariants, priors and interface design," in Human Haptic Perception: Basics and Applications, M. Grunwald, Ed. Basel: Birkhäuser Basel, 2008, pp. 381–392.

[16] A. E. Saddik, M. Orozco, M. Eid, and J. Cha, "Haptics: General Principles," in Haptics Technologies: Bringing Touch to Multimedia, A. El Saddik, M. Orozco, M. Eid, and J. Cha, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 1–20.

[17] H. Seichter, R. Grasset, J. Looser, and M. Billinghurst, "Multitouch interaction for Tangible User Interfaces," in 2009 8th IEEE International Symposium on Mixed and Augmented Reality, 2009, pp. 213–214.

[18] M. Konyo, S. Tadokoro, A. Yoshida, and N. Saiwaki, "A tactile synthesis method using multiple frequency vibrations for representing virtual touch," in 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005, pp. 3965–3971.

[19] F. Shao, T. H. C. Childs, C. J. Barnes, and B. Henson, "Finite element simulations of static and sliding contact between a human fingertip and textured surfaces," Tribol. Int., vol. 43, no. 12, pp. 2308–2316, Dec. 2010.

[20] W. H. Talbot, I. Darian-Smith, H. H. Kornhuber, and V. B. Mountcastle, “The sense of flutter-vibration: comparison of the human capacity with response patterns of mechanoreceptive afferents from the monkey hand.,” J. Neurophysiol., vol. 31, no. 2, pp. 301–334, Mar. 1968.

[21] V. G. Macefield, C. Häger-Ross, and R. S. Johansson, “Control of grip force during restraint of an object held between finger and thumb: responses of cutaneous afferents from the digits,” Exp. Brain Res., vol. 108, no. 1, pp. 155–171, Feb. 1996.

[22] “DINED.” [Online]. Available: https://dined.io.tudelft.nl/en. [Accessed: 14-Aug-2019].

[23] R. Ackerley, I. Carlsson, H. Wester, H. Olausson, and H. Backlund Wasling, “Touch perceptions across skin sites: differences between sensitivity, direction discrimination and pleasantness,” Front. Behav. Neurosci., vol. 8, Feb. 2014.

[24] M. Jones and G. Marsden, Mobile Interaction Design. USA: John Wiley & Sons, Inc., 2006.

[25] R. T. Verrillo, “Investigation of Some Parameters of the Cutaneous Threshold for Vibration,” J. Acoust. Soc. Am., vol. 34, no. 11, pp. 1768–1773, Nov. 1962.

[26] S. J. Bolanowski, G. A. Gescheider, R. T. Verrillo, and C. M. Checkosky, “Four channels mediate the mechanical aspects of touch,” J. Acoust. Soc. Am., vol. 84, no. 5, pp. 1680–1694, Nov. 1988.

[27] R. M. Peters, E. Hackeman, and D. Goldreich, “Diminutive Digits Discern Delicate Details: Fingertip Size and the Sex Difference in Tactile Spatial Acuity,” J. Neurosci., vol. 29, no. 50, pp. 15756–15761, Dec. 2009.

[28] R. S. Johansson, U. Landstro¨m, and R. Lundstro¨m, “Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements,” Brain Res., vol. 244, no. 1, pp. 17–25, Jul. 1982.

[29] T. A. Kern, Engineering Haptic Devices: A Beginner’s Guide for Engineers, 1st ed. Springer Publishing Company, Incorporated, 2009.

[30] J. Hartcher-O’Brien, B. Edin, and V. Hayward, “Shape-elasticity tactile confound, Hand, Brain and Technology: The Somatosensory System,” CSF Conf., Aug. 2018.

[31] N. Life, “Feature: Taking A Look Back At The Nintendo 64 Rumble Pak,” Nintendo Life, 11-Jul-2013. [Online]. Available: http://www.nintendolife.com/news/2013/07/feature_taking_a_look_back_at_the_nintendo_64_rumble_pak. [Accessed: 14-Aug-2019].

[32] S. Kim and G. Lee, “Haptic Feedback Design for a Virtual Button Along Force-displacement Curves,” in Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, New York, NY, USA, 2013, pp. 91–96.

[33] L. P. Kirsch et al., “Reading the mind in the touch: Neurophysiological specificity in the communication of emotions by touch,” Neuropsychologia, vol. 116, pp. 136–149, Jul. 2018.

[34] A. M. L. Kappers, “Human perception of shape from touch,” Philos. Trans. R. Soc. B Biol. Sci., vol. 366, no. 1581, pp. 3106–3114, Nov. 2011.

[35] “Coding of pleasant touch by unmyelinated afferents in humans | Nature Neuroscience.” [Online]. Available: https://www.nature.com/articles/nn.2312. [Accessed: 14-Aug-2019].

[36] R. C. Vortherms, “Clinically Improving COMMUNICATION Through Touch,” J. Gerontol. Nurs., vol. 17, no. 5, pp. 6–9, May 1991.

[37] C. Moussette, “Simple haptics : Sketching perspectives for the design of haptic interactions,” 2012.