SP3X: A Six-Degree of Freedom Device for Natural Model Creation

Richard Henry Whitney, III

Bachelor of Science in Biomedical Engineering
Johns Hopkins University, June 2004

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences at the
Massachusetts Institute of Technology
[February 2007]
January 2007

author

Richard H. Whitney, III
Program in Media Arts and Sciences

thesis supervisor

Hiroshi Ishii
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences

chair, departmental committee
on graduate studies

Dr. Andrew B. Lippman
Program in Media Arts and Sciences

SP3X: A Six-Degree of Freedom Device for Natural Model Creation

by

Richard Henry Whitney, III

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on February 2, 2007, in partial fulfillment of the
requirements for the degree of
Master of Science in Media Arts and Sciences

Abstract

This thesis presents a novel input device, called SP3X, for the creation of digital models in a semi-immersive environment. The goal of SP3X is to enable novice users to construct geometrically complex three-dimensional objects without extensive training or difficulty. SP3X extends the ideas of mixed reality and partial physical instantiation while building on the foundation of tangible interfaces. The design of the device reflects attention to human physiologic capabilities in manual precision, binocular vision, and reach. The design also considers cost and manufacturability. This thesis presents prior and contributing research from industry, biology, and interfaces in academia. A study investigates the usability of the device and finds that it is functional and easily learned, and identifies several areas for improvement. Finally, a Future Work section is provided to guide researchers pursuing this or similar interfaces.

The SP3X project is a result of extensive collaboration with Mahoro Anabuki, a visiting scientist from Canon Development Americas, and could not have been completed without his software or his insight.

Thesis Supervisor: Hiroshi Ishii
Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences

SP3X: A Six-Degree of Freedom Device for Natural Model Creation

Richard Henry Whitney, III

The following people served as readers for this thesis:

Thesis Reader____________________________________

V. Michael Bove, JR
Principal Research Scientist
MIT Media Laboratory

Thesis Reader____________________________________

William J. Mitchell
Professor of Architecture and Media Arts and Sciences
MIT Media Laboratory

Acknowledgements

This thesis is the result of efforts far beyond my own.

First and foremost I would like to thank my advisor, Hiroshi Ishii. His continuing faith in my abilities has made this thesis possible. Our three-hour discussions have made me a wiser person, and given me a deeper understanding of all things academic and tangible. I have been inspired by your work, and hope to one day be worthy of your admiration.

My interactions with my thesis readers were inspirational and invaluable. Michael Bove and William Mitchell have granted me boundless patience, far beyond what could ever be expected.

Jim, Angela, Jim, and Ari for the late nights, the advice, the friendship and the food.

Claude and Michael for their unflagging encouragement and for the opportunities they have given me.

All the past and present members of TMG for the inspiration, critique, and atmosphere that made this thesis happen.

Kelly and Ernesto for friendship through some odd times.

Joe, Andy, Mike, and the other 307 guys for their enduring friendship.

Big Willie, for always being there, from the very beginning.

Darth and Scott for being the go-to guys for anything, their enduring friendship, being endless fonts of motivation, lending their technical abilities, and for being smarter than I... at least in a couple of things. I cannot say enough about these two, but leave it at I could not have finished this without you. Thank you.

Special thanks to John and Tom for all their work maintaining the shop.

Linda Peterson and Pat Solakoff deserve recognition for their support, both officially and personally. Linda, Mountain Dew isn't an addiction – it's a lifestyle.

Lastly but not certainly least, I want to thank my family for their love and support. I wouldn't be here if it weren't for you.

To my beloved wife, with whom all things are possible.

Contents

Chapter 1

Introduction

For millennia, artists and designers have honed manual skills in order to better express and convey ideas. The talents of visual and physical communication set them apart and bestowed upon them the ability to realize their visions. In the modern world, the artisan skills have been supplanted by a tool that, while more limited than traditional media, could be considered no easier to learn: computer-aided design software. Computer-aided design has become a primary mode of expression out of necessity. For most commercial applications, designs must be digitized to enable their manufacture by standard mass-production technologies. The constraints of these manufacturing technologies ripple up the production process and impose restrictions on designers. These constraints are clearly visible in most computer-aided design packages, because they were intentionally incorporated to aid the commercialization pipeline. While these constraints are desirable from a consumer product design standpoint, they obstruct and encumber the expression of more natural and organic forms [40].

Recent trends in product design and culture herald a return to smooth, nature-inspired contours and press for individuality and customizability. These rounded forms are more visually appealing and comfortable than the hard edges of purely engineered shapes. The field of physical ergonomics anchors the movement with scientific support, and has promoted industry adoption [40].

1.1 Problem Description

The emergent problem is a lack of affordable, easily understood tools for creating natural 3D models. Designers, engineers, and artists need a new modeling tool that reduces the mathematical obfuscation of current 2D CAD approaches, and frees them to express the organic forms that are both visually appealing and physically comfortable. The proposed tool should be easy to learn and use in order to promote adoption, be relatively inexpensive, and confer 3D affordances that are helpful and not widely available.

The goal of this research is to design and implement a simple tool for artistic 3D model creation that attempts to leverage basic human skills and biological characteristics. The key human features we hope to incorporate are spatial perception, proprioception, bimanual input, and manual dexterity.

A perfect digitally enhanced clay presents an ideal solution to the proposed guidelines. Perfect digital clay represents a long-envisioned goal for designers of 3D modeling tools. This digitally enhanced clay retains all of the flexibility and tactility of real clay, but gains several critical features from the digital world [38]. The capability to save and load structures is one of the many desired additions to physical clay, as it opens up many new opportunities. Sculptures could be transmitted over the internet and recomposed on-site, without shipping. Save/load also makes perfectly geometric constructs easy, with mathematically defined files. One central feature is the capability to define the digital "undo" command in the physical world, yielding a consequence-free sculpting environment: any mistakes can vanish with a single command. These points led to a design that favors the virtual techniques over the complete physical structure of the digital clay concept.

1.2 Creating Flexible Physicality

The Spatial Planar 3D Extruder (SP3X) is a novel device that attempts to loosen the creators restraints, while maintaining a focus on implementation practicality. SP3X enables users to create a set of natural 3D objects through a human-scale tangible interface, coupled with a mixed-reality display. In effect, a SP3X sculptor can see the virtual creation in 3D physical space, at a 1:1 scale, and manipulate it via a simple two-handed interface that makes use of the remarkable powers of human manual dexterity and improvisation.

SP3X blurs the line between mixed reality and tangible interfaces by combining elements of each. The result can be classified as a "tangible controller" within a mixed reality environment.

This work has been the result of a highly productive collaboration with Mahoro Anabuki of Canon Development Americas, who provided not only the bulk of the software but also design feedback and productive discussion.

1.3 Document Overview

Chapter 2 provides a brief introduction to SP3X by means of an example of simple usage in a CAD scenario.

Chapter 3 provides the background necessary to frame SP3X in a larger context, defines significant physiologic terminology, and discusses the limitations of clay. Finally, it summarizes a series of projects that are related and sources of design inspiration.

Chapter 4 describes the theories underlying the design of the SP3X project, as they extend from the works of the previous chapter. The chapter goes on to chronicle the implementation of the device, and concludes with concerns and limitations.

Chapter 5 details a series of user studies which were conducted to evaluate the device, and the results of these experiments. Selected works of the participants are presented.

To finalize, Chapter 6 concludes the thesis document with recommendations of future changes, alternative designs, and possible applications of the SP3X system.

Chapter 2

Example

SP3X works by capturing the profile of a physical loop on the interaction surface. The loop is comprised of twenty segments that are strung like the beads on a necklace. The beads have fixed distance spacing and bind magnetically to the working surface. This prevents them from being moved unintentionally, as when the surface becomes steeply tilted. The strength of the magnets is low enough that they are still easily moved by hand.

The interaction surface, or "deck", translates manually through 3D space, with 6 degrees of freedom rotational (roll, pitch, yaw) and translational (x, y, z). The SP3X system tracks the location of the deck and the resident loop.

The following example is intended to introduce the SP3X system via a typical usage scenario.

2.1 A Basic Interaction Scenario

The user dons a lightweight head-mounted display (HMD) and becomes immersed not in virtual reality, but in the real world. The Canon mixed reality headset captures video via a pair of cameras that are approximately in-line with the users eyes, and redisplays these video streams to the user on two small LCDs, one for each eye.

The user sees a single virtual element in their otherwise unaltered (albeit lower resolution) surroundings. This element represents the digital capture of the physical loop, or "slice", and acts as a preview of the current outline. It also functions as a confirmation that all of the segments have been accurately acquired.

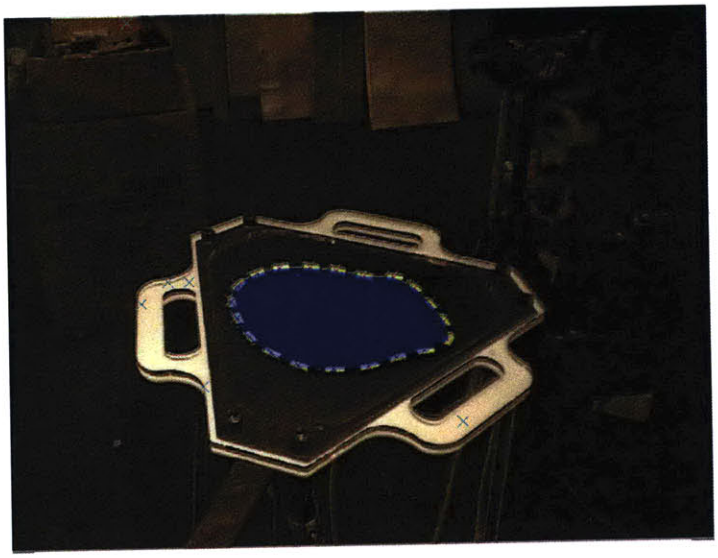

The loop deforms easily when the user acts on it. The user can pinch, poke, press, impress, or stretch the loop, or any of a myriad of other basic physical behaviors from simple to complex. These behaviors can be applied in tandem or series to sculpt the loop into the desired shape.

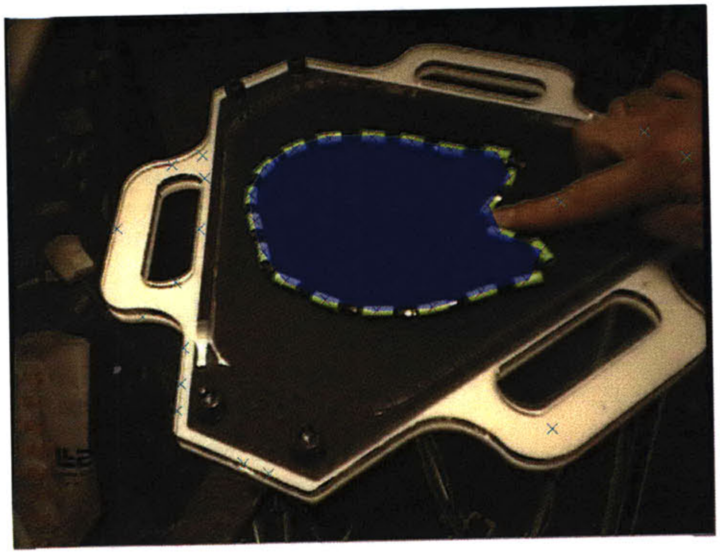

As the user moves to look around the environment, the representation of the virtual slice changes with the perspective of the physical deck, exactly as a sheet of paper lying on the deck would. The virtual slices location is affixed to the deck, and this permits the user to inspect it and change viewpoints in the most natural way - by moving and looking, just like any real object.

This natural approach to 3D space navigation is both more intuitive than conventional CAD navigation due to its identical matching to the real world analogue, and more efficient because navigation and modification can occur simultaneously (spatially multiplexed).

By depressing a foot pedal, the user captures or "finalizes" the active slice. The slice becomes fixed in shape and location, and the users actions on the loop no longer alter this slice. Upon slice finalization, the system changes the display color of the old slice from blue to red to indicate that it has been captured and is now passive. At the same time, a new slice automatically appears and becomes the target of loop modifications.

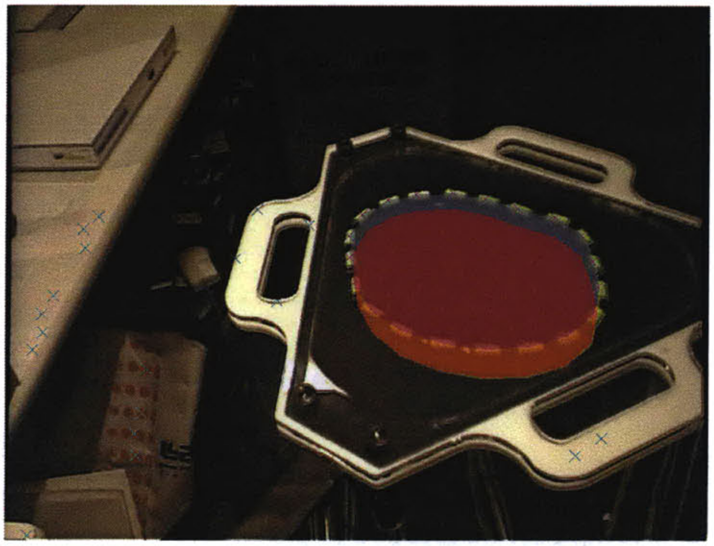

If the user grasps the handles of the deck and depresses a switch, the normally rigid supporting substructure suddenly goes completely limp. The deck can then freely move to another location with few constraints and very little resistance. Upon arriving at the desired locus and orientation, the user simply releases the switch and the deck returns to its solid prior state.

When two or more slices exist (n-1 passives and 1 active) these slices are

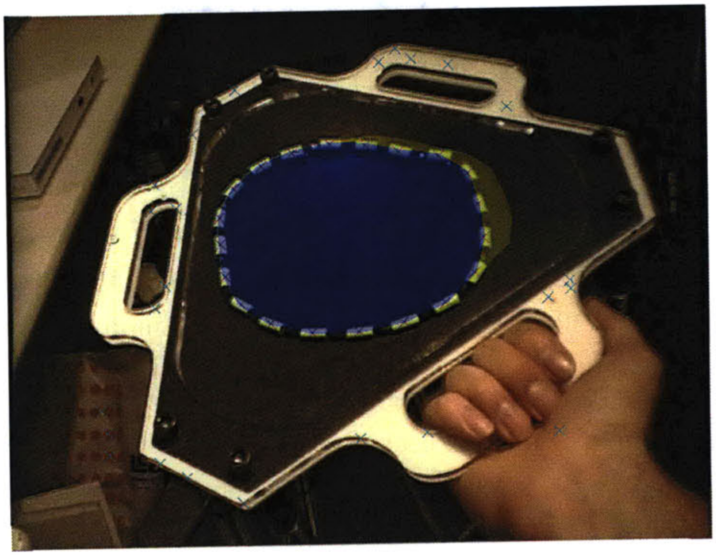

geometrically connected and rendered to form a digital object. The segment connecting the active and most recent passive changes to reflect any alteration in the outline or location of the active. This continuous preview maintains the users awareness of the virtual components of the object being created.

Complete models are built up layer by layer in this manner. However, a parallel approach is also available. In the continuous extrusion mode, layers are added at regular intervals as long as the pedal remains depressed. This mode is helpful if the user wishes to create a long pipe form without altering the cross section.

Chapter 3

Background and Related Research

3.1 Background

3.1.1 Manufacturing and Design

In the beginning, humans modified the world around them and introduced the first artificial structures through intent and design. They created clothing and tools by hand from the resources available, and every item was unique. The items reflected both the manufacturing process and the materials: irregular and unique. The emergence of quality-controlled modern manufacturing and homogenous materials eliminated the variability and character of consumer products, condemning them to sameness.

Consumer product design and manufacturing technologies have driven each other throughout human history. They are intimately linked in the process of product creation. An object cannot be built until it has been envisioned, and a vision cannot be embodied until the required production techniques exist.

Consumer product design has always been required to adapt to cost constraints in order to yield products that are affordable to the targeted segment of the population. Depending on the era, the components of manufacturing cost can vary. While materials represent a base cost, assembly, skilled labor, and tooling[1] are dependent on the technical process. The need to meet price windows has, to a large extent, driven design.

Early design was free of the constraints imposed by machining and mass production. All items were hand-made and thus there were no tooling limitations, planarity requirements, or issues with the draft angle of a die. Three factors dictated the total expense of a product: materials cost, time required, and degree of skill necessary. Each piece was unique, a one-off this requires significant man-hours of labor per part, and so the total cost was high. Intricate or ornate items took the skills and time of an experienced craftsman, and pushed the cost even higher.

The costs of producing a particular object are also heavily dependent on the manufacturing technologies employed. When new cheaper and more versatile technologies emerge, designers gain flexibility in what they can create within their cost constraints.

Mass Production Technologies

The emergence of mature mass production technologies at the end of the 19th century brought a new cost structure and with it, new design constraints. In the mass production structure, a significant portion of the cost of a product is incurred during setup. The setup includes the creation of molds, development of workflow, and tool configuration. After the setup is complete, the additional cost of each item is negligible and attributable almost entirely to the raw materials. However, changing any element of the product cascades through the workflow and quickly becomes expensive. While designers could reach millions of consumers with a successful product, they were now limited by what the machines could readily produce.

DFM and DFA

Design for manufacture (DFM) and design for assembly (DFA) are two schools of thought for reducing production costs. Each imposes guidelines on product design from the earliest stages. DFM focuses on em-

---

[1] Tooling is the planning, creating, and arranging tools for a particular manufacturing process

ploying only standard components; including draft angles on parts; and more than a dozen other principles to reduce cost and difficulty in the use of conventional manufacturing technologies. DFA adds more constraints like "reducing parts count" and "reducing handling time". Between these two sets of guidelines, designers must incorporate so many factors that their output is dramatically limited from the very beginning. Moreover, the financial consequences of ignoring these principles are economically unthinkable.[17]

Rapid Manufacturing

The latest step in the evolution of production technologies is referred to by the umbrella term rapid manufacturing. The unifying principle amongst these technologies is the capability of creating parts and/or complete products without expensive tooling. With the elimination of tooling and work-up costs, the initial cost of manufacturing an item is virtually eliminated. This makes it possible to create very small runs, or even single item runs, without incurring dramatic financial penalties due to the distribution of the initial cost.

Rapid manufacturing is comprised of two varieties- additive and subtractive. Subtractive technologies, as the name implies, involve the subtraction of material from a solid mass of starting media. Laser cutting, abrasive water jets, and wire EDM remove material from a plate of material in a two-dimensional path, leaving a flat cut-out. This approach is well-suited to many applications like gears and fixtures, and multiple two-dimensional pieces can be combined into three-dimensional structures.[17]

Additive technologies, conversely, begin with an empty location and add very small amounts of material that bind to each other and build up. Additive systems can create complicated structures at the sub-millimeter scale, and even completely sealed monolithic mechanical systems- two feats that are impossible with conventional manufacturing technologies. Additives can employ a variety of materials, and some systems can use multiple materials within the same object. For example, a modern selective laser sintering machine (SLS) can form a hard plastic toothbrush handle with soft rubber grips as a single unit, without other tools or steps. Other additive technologies include plaster and starch 3D printers and fused deposition modeling (FDM), which builds objects of ABS plastic. [17]

Rapid manufacturing technologies enable completely individualized products and help to bring about the advent of mass customization.

3.1.2 Why Not (Physical) Clay?

The quest for the ideal digital clay, or approximation thereof, has been a persistent feature of the human-computer interaction community since its origin. This quest has been laborious and expensive. All this effort begets a simple question: why not clay? Wouldnt it be easier to just build the models of out of clay? The answer is multifold.

Need for a Digital Representation

A practiced artist can create beautiful and complex flowing forms from a shapeless lump of clay. The artist can sculpt a perfect model of a future product, but the clay model must then be converted to digital information to be useful to the modern product cycle. Without the digital representation, tools cannot be programmed, dies cannot be milled, and packaging cannot be designed. Moreover, consumer design is often geographically distributed; it is far faster, safer, and economical to transmit files than to ship clay.

In future of mass customization, the consumer will design the purchase and manufacturing will transform the customers design into the product. The notion that a consumer will buy clay, physically sculpt the object, fire it in a kiln to prevent in-transit deformation, ship it to a manufacturer, and then receive the result weeks later is theoretically and technically possible, although inefficient and impractical.

It is possible to bridge the clay to digital chasm through 3D scanning technologies. Expensive commercially available devices like non-contact 3D laser scanners, which can cost upwards of $40,000, map the surface of an object using thousands of sample points. The resulting "point cloud" can then be interpreted into a digital representation of the original object via a combination of computer and human effort.

This scanning technology carries some significant shortcomings. First, the technique has a minimum feature size- the smallest element that the scanner can "see." Below this threshold, all information is lost. Second, the scanner only scans surfaces that are visible from a specific viewpoint. The object may be rotated around the vertical axis on a turntable to reveal its other sides, but internal structures and occlusive concavities cannot be captured. Third, the device is well beyond the price range of the vast majority of individuals. More affordable 3-D scanners are limited in abilities and resolution, just exacerbating the first two barriers. These vital limitations put the scanner approach out of the mass customization vision.

The technical capability of physical to digital conversion does not lessen the significance of the other disadvantages of clay.

1:1 Scale

Full-scale development is widely considered to confer great advantage. The impact of most product designs is heavily visual, and this impact becomes diminished when the product prototype is seen at 1:10 scale. A two-dimensional image of the prototype detracts even more from the visceral presence that a real, full-size object exudes.

Automotive companies in particular spend a great deal of money and time on the creation of 1:1 scale models. A team of sculptors craft these models from clay, but the work is not original. The team tries to perfectly recreate a digital model created by a designer using some variety CAD. The clay model is therefore redundant information. The model is produced, though, because it has great value to the designer. The designer needs to have a real spatial and visceral perception of the model. This power cannot be conveyed by a two-dimensional image (screen or paper) or a small three-dimensional model.[47]

Aside from the expense and inefficiency of entirely recreating the model to a 1:1 scale, an even more significant problem with scale lurks: perspective. While the top level designer may be most concerned with the contour of the primary strokes along the profile of the car, which let us say is most efficiently designed at a 1:50, the design of a crevice may necessitate 2:1 perspective. While this can be obtained through CAD or clay alone, neither allows the freedom to move in 3-D between the scale used for design and 1:1, which is necessary to put the overall design in perspective.[21]

Clay faces limitations of scale, which cannot create equal impact and be overcome by digital 2-D visualizations. In addition to this innate and unavoidable conflict between the nature of clay and the necessity of an effective medium, the expense of the materials, time and physical space inhibits the full utilization of the abilities of scaling.

Numerical Requirements

In the product design arena, there are often quantified dimensional requirements. Some components, like the screen in a cellular phone, may be determined before the housing is designed. Other elements like power connections and headphone jacks are of standardized dimensions. These place a burden on the designer to match the sizing exactly or find themselves with incompatible components and a product that cannot be assembled.

The application of these physical restrains to clay can be very difficult, as it is difficult to set precise parameters while sculpting the design. The consequences of mistaken or even slightly incorrect dimensions can alter the original intent of the design, as it may have to be stretched and skewed to match the quantified dimensional requirements.

Requires Skilled Individual

A critical shortcoming of clay lies in the necessity of an artist. While design requires some degree of artistry, the ability to accurately sculpt a design requires training which often must be taught, developed, and refined over many years. This makes it difficult if not impossible for an amateur, or customer, to design precisely what he or she envisions. This "high floor, high ceiling" medium is difficult to use and master, despite the breadth of potential to the skilled artist, making it practically inaccessible to the common user. If free-form design can only be conveyed through a highly skilled talent, then we have no hope for the mass-customization future. The ideal digital future will develop a process by which anyone, amateur or professional alike, without limiting the applications potential.

Malleability

During the sculpting process, every phase before the final structure is lost. It requires the sculptor to manually revert to former incarnations of the design, rather than having such information and previous phases saved. This is even more significant when the clay is used in an incremental or recursive design cycle. In each of these design cycles, a feedback loop requires numerous revisions, and thus in clay design, many iterations become costly, time consumer, redundant and difficult to precisely recreate.

Once clay has entered the kiln, it is no longer malleable. In order to then alter a feature on the prototype design, the creator must choose whether to recreate the sculpture or rely on less intuitive computer applications. The data is thus difficult to post-process, making revision and reverting to an earlier phase difficult. Ideal clay would allow the designer to interact with the model and start from the last revision, or revert to a previous revision.

Conclusion

Ultimately, the future of product design and mass customization is incompatible with the physical medium, clay. Clay fails to adhere to the ideals and necessities of product design.

3.2 Related Research: Physiologic

Human physiology plays a large role in the design of any human-computer interface, whether it be the limited hearing range of the ear or the resolution and motion sensitivity of the eye. In the case of SP3X, human tactile and mechanical parameters were of key importance. Muscular control precision at the joints in the arm; proprioception and kinesthesia; and average human factors measurements (strength, height, and reach) were core scientific considerations and supporting elements.

3.2.1 Proprioception and Kinesthesia

Proprioception is one of the bodys lesser-known senses. This is the sense that allows you to touch your nose without the use of your sight, as in a completely dark room. Proprioception falls in a category named interoception. These are the senses that your body uses to monitor its position and orientation. This category also contains balance. The other general sensory category is exteroception, and this is composed of what are commonly thought of as the five senses- touch, sight, taste, sound, and smell [48].

Kinesthesia is the "memory" of the muscles. This allows you to recall motions and locations in space. It could be considered the memory of a single or series of proprioceptive states. Though the two terms are commonly used interchangeably, the roots of these words show that to be incorrect. Kinesthesia also makes it possible to become better at a task by repeating it.

3.2.2 Manual Precision

The precision of control that humans can exert varies both by individual and by joint. While the variability between individuals is beyond the scope of this endeavor, the relative precision of the joints is certainly not. An interface that takes these differences into account can present a control structure that leverages the strengths to improve the detail of control.

Humans have much greater control in the hands and wrists than they do in the elbow and shoulder. This comes from two primary physiologic factors. These are the relative lengths of the structures that the joints control and the amount of representation in the brain. Together these strongly favor precise manual dexterity and coarse shoulder power [6].

The first is the relative lengths of the actors. The forearm is longer than the hand (generally by about two-thirds), and therefore an error of 1 degree at the elbow produces a larger translational error than the same 1 degree at the wrist.

The second is the size of representation in the brain. The size of a region responsible for a body part is a direct indicator of how precise the control over that part is. This was thoroughly tested by Rioult-Pedotti, et al in their seminal work on the development of greater manual precision in rats [34]. The rats were trained to use their paws more heavily by placing food pellets inside a box with one small entrance hole. The rats had to reach inside the box and grasp the pellets. After five days of training, their brains were analyzed and it was determined that the area occupied by the forepaw had increased in size relative to the hindpaw, which had remained the same. In humans (as in most mammals who have them), the hands are a much larger area in both sensor and motor layers of the brain.

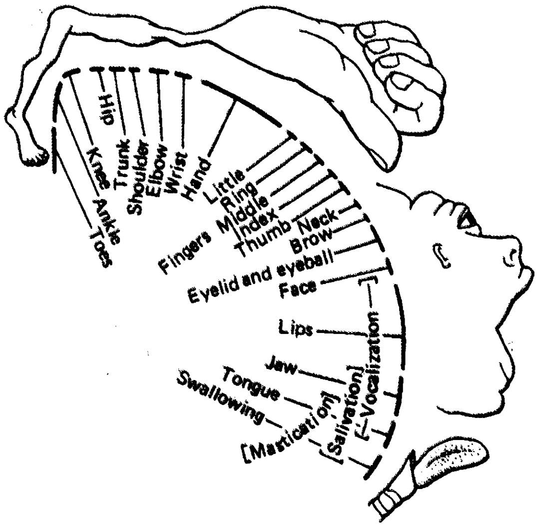

Figure 3-1: A motor homunculus, indicating the representation of body regions on the brain.[6]

3.3 Related Research: Interfaces

The related research is extremely broad. SP3X has been inspired and influenced by the designs and philosophies of many more projects than can reasonably be described within the bounds of a thesis. We have created a new system categorization that can characterize the salient features of the relevant interface research. This system relies on two axes: physical bandwidth and layer of physical/virtual division.

3.3.1 Bandwidth

In this context, the term "physical bandwidth" has little relation to its use in electronics and software. Communication rates from the human to the computer cannot be expressed as bits per second. The distinctions are far more dramatic and difficult to quantify.

Physical bandwidth is an observational measurement of an interfaces capability to receive input from human manipulators in parallel. An interface that allows the user to press two buttons simultaneously, therefore, possesses greater physical bandwidth than one which senses a single button.

The physical bandwidth scale is defined from the users perspective. It can be thought of as spanning the space from no tactile input (watching a movie in a theater) to the full tactile capability (playing in sand).

3.3.2 Layer of Division

Human-computer interfaces are designed to convey user input into the digital domain. In the case of physical input, every interface uses the same generalized pathway to pass this information. Users manipulate tools, tools modify objects, and objects tie into the rest of the digital world. In a grammatical sense, tools are verbs (sculpt, hammer, drill) and objects are nouns (clay, nail, wood). Tools perform a very specialized function, and can be physical or virtual. Objects contain data and represent the digital medium.

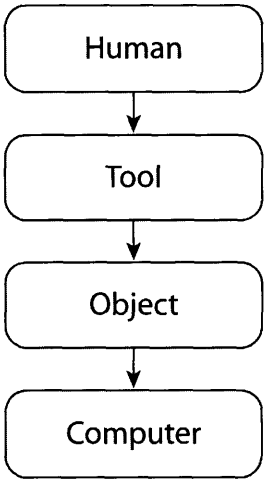

Layer of division between the user and the virtual world.

3.3.3 Summary

In the interests of feasibility, all human-computer interfaces must be designed as a series of compromises. These decisions can reflect interface goals, technical limitations, or financial constraints. The best input devices are designed to complement the data on which they act. Some systems use these input constraints beneficially.

3.3.4 GUICAD

Software applications have historically been two dimensional, and the associated interface devices have mirrored this planarity. The staple tool of graphical user interfaces (GUIs), the mouse, is a perfect example. The mouses design and operation rely on contact with a surface to perform, so the working range of motion and control are necessarily within the planar physical constraint.

The majority of todays user interfaces continue to be two dimensional. The mouse remains a principle GUI device, supplemented by the character input of the keyboard. This planarity has come about due to a variety of factors, history and cost being significant contributors. The problem arises when these inherently two dimensional controls (and displays) are applied to three dimensional application spaces. Three-dimensional scenes are downsampled to 2D, losing depth data and intuitive spatial navigation. Then, 2D interaction is upsampled back to 3D by sequential input. The information and time sacrificed by these "lossy" transformations is massive.

GUI CAD was expressly designed for and evolved to support the modeling of components for manufacture (physical output). The virtual tools are numerous and varied, and all exist to aid in specifying parts. Some of the virtual tools are directly mapped to physical counterparts, like the fillet function and a router.

GUI CAD has very limited physical bandwidth, and is completely virtual. The combination of constraints helps to keep users within the guidelines of readily manufactured shapes. If an engineer was presented with a completely free-form interface, like a lump of clay, s/he would probably have a great deal of difficulty molding parts with the precision tolerances necessary for mechanical operation.

3.3.5 Mixed Reality

The central tenant of mixed reality (MR) proposes that elements of the real and virtual worlds should be displayed together. Mixed reality encompasses the entire spectrum of real and virtual world combinations. This breadth includes both the bringing of physical objects into the virtual world, and virtual objects into the physical world.

In the case of SP3X, we are most concerned with the augmented reality end of the MR spectrum. Augmented reality systems overlay virtual information and constructs on the physical world. This virtual information is anchored to objects or locations, preserving the illusion that the user interacts with an unaltered reality. The use of head-mounted displays (HMDs) is a common approach, as is the use of digital projectors. Six degree-of-freedom tracked HMDs provide an excellent solution for applications with 3D virtual objects, as they provide for true 3D inspection. Natural head motions replace artificial device-based world navigation, making the interaction more intuitive and seamless.

A commonly applied implementation of MR is the ARToolkit [19]. This system, like several others [35], makes use of visual tags. These visual tags are images, generally black and white, and have a definite orientation and can be readily processed by a computer. These tags function as the real-world anchor for virtual objects. Through the combination of HMDs or other camera-connected displays and the visual tags, creative effects can be produced. An example would be a virtual potted plant that always appears to rest on a physical domino. The result is an interface with a high degree of physical/virtual coincidence, but no physicality.

Mixed reality maintains strong philosophical differences from tangible user interfaces.

3.3.6 Tangible User Interfaces

In contrast to mixed reality interfaces, tangible user interfaces (TUIs) are defined by the physical representation of their components. Objects in TUIs commonly possess a form that describes their representation or function. The Amphibian [8] project allowed users to associate digital information with physical objects that they found to be visually descriptive. This approach stands sharply in contrast to the cryptic and planar codes of the AR Toolkit.

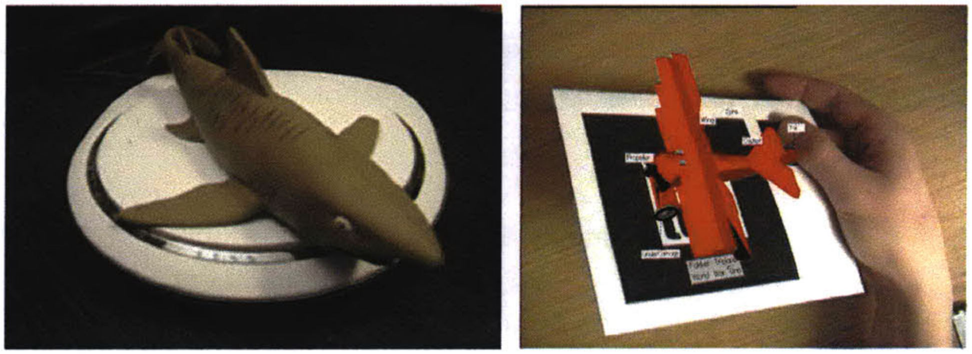

Figure 3-2: Shark representative object from Amphibian[8] and ARToolKit tag[19].

Ishii proposed tangible user interfaces in the paper "Tangible Bits" [18], pressing for a new interface philosophy to replace the temporally multiplexed structure of GUIs. Ishii emphasized the need to "embrace the richness of human senses" as well as making use of the "skills people have developed through a lifetime of interaction with the physical world" [18].

Tangible user interfaces have penetrated deep into the human-centric aspect of interaction. By coupling controls and data to physical objects, TUIs have disambiguated the dynamic and contextual function remapping of the mouse/pointer modality, and improved the human readability of applications. The physicality of tactile interfaces encourages exploration by experimentation [32], and continuous feedback supports this endeavor. Users are kept apprised of virtual control status by the physical state of the objects. In the context of tabletop systems, TUIs generally maintain a 1:1 mapping of physical objects and virtual representation or function, and this leads to a proliferation of graspable tokens on the work surface. Far from being a detriment, the availability of many simultaneous interaction points, makes complex bimanual behavior possible. The Bricks [10] study found that, in the simple sorting of LEGO blocks, subjects tended to use both hands in complicated continuous choreography, complimented by multi-object group gestures.

Tabletop TUIs

Tabletop systems emerged early in the development of tangible user interfaces due to a readily understood metaphor (the desktop) and the relative ease of sensing and displaying in two dimensions. Tabletop TUIs evolved rapidly from two object systems like GraspDraw [10] into multi-puck, dedicated purpose systems like the Sensetable [31] platform.

URP

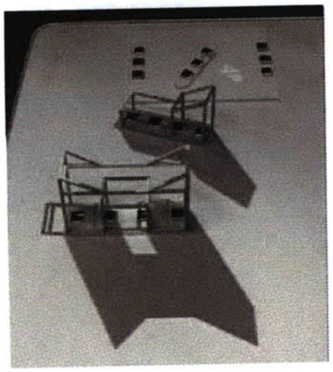

Underkoffler's URP interface.[45]

URP [45] is an excellent example of the conceptual flexibility of tabletop TUIs. The project consists of physical wireframe building models, two dials, and an oblong widget. The physical buildings directly represent their virtual analogues, even in structure. These wireframe models cast complete, solid digital shadows on the table, as though they were illuminated by an invisible sun. A digital wind swirls around the buildings, showing turbulences and wind tunnels.

One of the dials resembles a clock, but serves the reverse purpose- the clocks hand can be used to control the time of day, rather than display it. The other dial holds a similar inversion: despite looking like a weather vane, the dial controls the direction wind in the simulation. Unlike the dials, which alter the simulation at large, the widget acts on buildings. By tapping the "G" end against a building, the buildings simulated material changes to glass and begins to reflect light. The "B" end returns the material to the default state- non-reflective brick.

URPs significance lies in the physical representation of both tools and objects. The buildings work as object icons [44]. In our experience, users interact with the buildings first, because they are instantly understood. The physical structure of the wireframe model informs the user of the virtual representation.

With time and experimentation, people can ascertain the functions of the tools. The dials are relatively self-explanatory because they mirror familiar items. After a brief period of playful exploration, the behavior of the simulation can be correlated to the users pre-existing knowledge of the items.

The material widget, however, has no real world analogue and stands out sharply by comparison. Without instruction, few (if any) users discover the "tapping" operation that activates this tool.

PICO

PICO [32] is the most recent development in tabletop TUIs. PICO focuses on supporting computer-human negotiation through experimentation. In the primary application, users place and move pucks representing cellphone towers around a digitally projected map. Users try to maximize signal coverage by positioning towers and changing transmission power and lobe formation.


The PICO platform.[32]

The computer also moves the pucks around, in tandem with the user. The simulation governs puck motion based on two types of information. One type consists of geographic data, like zoning laws. The other is an optimizing algorithm that continuously moves the pucks towards a more efficient solution by small increments.

PICOs physical manifestation of simulation computation permits realtime human intervention. The user effectively "sculpts" the simulation along a desirable path. The physical nature of the interface allows the user to improvise constraints and barriers with available materials and custom tools.

Illuminating Clay

Illuminating Clay [33] was an investigation into continuous tangible interfaces. The user(s) sculpts a clay surface into a "2.5" dimensional topography. The system captures the sculpted surface, approximately in realtime, and computes a desired function on the data- options include local slope, shadow casting, and contour maps. The result is represented graphically, and displayed back onto the clay by a digital projector.

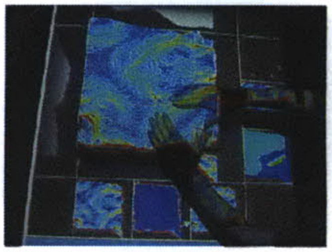
Piper's Illuminating Clay[33].

The system claims 2.5D rather than 3D because the surface is captured by a laser scanner suspended above the clay. This configuration makes horizontal concavity, like caves, impossible to detect. Despite this minor constraint, users respond very positively to the interaction, due largely to the intuitive nature of sculpting. A user can apply any combination or sequence of fingers, hands, or objects to the surface to achieve the desired structure. This improvisational capability encourages users to engage their natural creative behaviors.

Illuminating Clay is physically manifested at the object layer. The high degree of physicality (and correspondingly low degree of virtuality) is primarily responsible for the intuitive nature of the interface. The device behaves almost exactly like the user expects (sans concavity), so there is no learning curve. The trade-off is a lack of virtual flexibility, like loading and saving topographies.

3.3.7 Free-Form 3D Tools

Free-form interfaces often eschew physicality in order to make their flexibility technologically and financially feasible.

3-Draw

3-Draw [36] was one of the earliest interfaces to explore free-form data input. The 3-Draw system consisted of two six degree of freedom trackers and a standard monitor. Attached to one of the trackers, the "palette", was a flat rectangular piece of material, and the other tracker, used as a stylus, was instrumented with a button. The palette worked as the frame of reference for the virtual space, and controlled the orientation of the virtual model. The stylus drew curves through 3D space to define the model. The primary application space of 3-Draw was engineering.

3D Tractus

We encountered the 3D Tractus [22] project at a point when we were already far along in the design of SP3X. While we were initially surprised by the apparent similarity to SP3X, it soon became clear that the projects were guided by very different philosophies.

3D Tractus creates splines in a three-dimensional volume. The user draws on a tablet computer with a stylus. The tablet attaches to a framework that permits vertical motion. A linear encoder senses the vertical translation of the tablet, and the application combines this Z axis component with the X and Y from the tablet to give each vertex a 3D locus.

The 3D Tractus project contains some elements similar to that of SP3X. The Tractus device consists of a tablet PC mounted on a platform that can translate vertically. The vertical position and movement of the stylus are combined to produce a 3D spline. The resulting spline is displayed on the tablet PC in a top semi-coincident view, and in a smaller window on the PC from a side perspective view. The Tractus gains relevance to SP3X through its use of a moving frame of reference. The creators of the Tractus state that the novelty of their approach lies in the low cost of the device, but we believe that they overlook the human advantages of both a physical structure on which to work (to reduce fatigue) and a physical workspace to provide a tactile reference area that the user can tie into his/her mental model of the area.

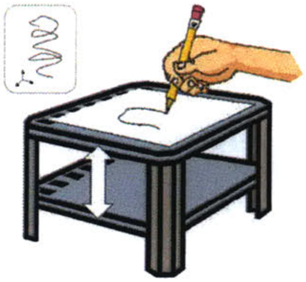

Basic operation of the 3D Tractus[22].

Surface Drawing

Schkolne, et al [37] observed that many artists had difficulty conceptualizing with 3D CAD software. They sought to combine pencil sketching and 3D modeling into a single interface that allowed artists to express themselves with free-flowing organic shapes.

Surface Drawing consists of a glove that is magnetically tracked in 6 DOF and two tools: a pair of tongs, a "magnet" and an eraser. When the user depresses a thumb switch, the curvature of the glove is captured as a continuous surface extrusion. These extrusions can be manipulated by the two tools. The tongs are used to move the drawing, and two tongs can be used in conjunction to scale the drawing. The eraser removes a very localized area of a drawing, and the magnet tool creates a local deformation.

Surface Drawing a curve through space[37].

The researchers noted that fully immersive environments were "too virtual" and detracted from the artists focus on the object. With this in mind, they employed a form of mixed reality using head tracking and a stereoscopic display referred to as the Responsive Workbench.

Subjects participating in a user study gave generally positive opinions, with only minor confusion as to the usage of the magnet tool. The interaction was intuitive enough for novices to produce recognizable objects within their first 20 minutes of exploration. During three months of weekly sessions, the artist Jen Grey created an artwork which merited display at the SIGGRAPH 2000 art show.

One interesting observation came from the researchers collaboration with the company Designworks. In the course of their exploration of how gesture and motion affect design aesthetics, they discovered that the designers could visually separate lines drawn "from the elbow" and those drawn "from the shoulder".

VRAD

In the paper "Smart Tools for Virtual Reality Based CAD", Fiorentino, et. al, [9] identify several advantages of virtual reality based design tools (VRAD) over traditional 2D CAD software. These are:

stereo vision: better shape perception (solids, free form);

spatial input: 3D interaction advantages from real life skills;

six degrees of freedom(6DOF): translation and rotation based control;

head tracking: enhancement of navigation and model insight;

mixed reality: fusion of digital (virtual) images with real world;

tangible interfaces: improvement of precisiondue to real object devices;

interaction freedom: encouragement of creativity and of the development of new ideas."

This list provides a succinct summary of the potential contributions of VRAD systems. The advantage descriptions closely follow the features of a desirable system laid out in the introduction.

The authors also isolate several human factors considerations of a VRAD interface. These constitute broadly known problems within the field, and are important to the real world usability of interfaces. The considerations listed are "pointing precision, fatigue, hand vibrations, lack of limb support, and interaction anisotropy".

The fatigue, hand vibrations, and lack of limb support concerns all stem from the use of an ungrounded, purely virtual environment where the user holds a pointing/interacting device. In this configuration the user moves the device through free space, as though s/he were gesturing with a wand or laser pointer. The user must support the weight of an object and his/her arm, often at full extension, for prolonged periods. Clearly, this can exhaust an untrained user quickly.

Hand vibrations extend from the same support issue. Small muscle twitches in the arm and shoulder, even breathing and heartbeat, produce a continuous string of movements. These motions are magnified by the length of the arm and culminate in oscillations of the hand approximating the desired position. This effect can be readily observed by aiming a laser pointer at a wall while holding ones arm at full extension.

Pointing precision can partially be attributed to these same appendage motion issues. The remainder of the error arises from working within a purely virtual environment. Without a concrete frame of reference with which to compare object size, it is difficult for users to accurately gauge the depth of a virtual object. [9]

In "Moving Objects in Space", Mine and Sequin [30] identify several other factors contributing to the weakness of virtual environment interactions, present further recommendations for VR interface design. They advise working within arms reach to improve control and to take advantage of Proprioception for object awareness. They cite the lack of haptic feedback in VR systems as a primary failing, and encourage the development of systems with physical working surfaces.

Chapter 4

Design and Implementation

4.1 Theory

The Spatial Planar 3D eXtruder (SP3X) was conceived with human physiology and the existing interfaces in mind. The design strives to incorporate the philosophies of tangible bits and mixed reality to create a novel device and interaction that embodies a compelling combination of technology and physicality.

Returning to the concepts of Physical Bandwidth and Layer of Division presented in the previous section, we sought to maximize the physical engagement and manipulation flexibility of the SP3X system. We were able to compare the contributions of different interface structures by ranking them in this metric, and thus to choose the configuration which we thought best met our objectives.

Layer of Division is the boundary at which an interface crosses from physical to virtual. This division occurs at one of the transitions in the User - Tool - Object - Computer chain. Within the context of the SP3X project, the Object is our digital version of clay and the Tools are anything used to manipulate it.

We considered the use of a physical handle with a purely virtual tool tip which could be changed by selecting different modes. Along the same path, we discussed the possibility of offering a full set of familiar physical tools, each representing a function (carving, adding, smoothing). The second idea holds more closely with the philosophies of tangible interfaces by being spatially multiplexed, but neither solution was appealing.

The root problem stemmed from the completely virtual nature of the object. The layer of division was too shallow, and this restricted user interaction to the predefined tool objects. With this constraint, we would need to create a vast array of these tools to provide the user with the desired level of flexibility.

The alternative was to make more of the interface physical and reduce the layer of division. In doing so, we would make the digital clay physical. The challenges that a perfect physical digital clay impose push such a concept into the range of impossibility. The basic requirements would include micro-scale actuation, tiny granule size, millions of sensors, and ad hoc networking. Therefore, we chose to assault the old problem with a novel approach.

The central philosophical premise of SP3X is that physicality is defined by perception. If an object is visible but out of reach, we assume that it is real. If we grasp a long pipe, we assume that the entirety of the pipe is real, despite having touched only a small part. Extending this philosophy to interface design, only the elements that the user can touch at any given point in time need to be physical. Whenever an element is not being actively touched, that element can be purely virtual.

Therefore, we propose an interface which relies on partial physical instantiation. The interface allows the user to directly manually manipulate the entire virtual world, but at any instant constrains this interaction to a single plane of physical handles. We believe this to be the most financially and technologically efficient method of creating a free-form sculpting device.

Partial physical instantiation can be classified as a subset of hybrid interfaces. Hybrid interfaces, in the field of mixed reality, incorporate elements of both the physical and virtual worlds in a seamless fashion. Users interact with physical elements that appear continuous or connected to other elements, but the untouched elements are purely virtual, substituted into the environment by the clever use of 3D display technologies. [20]

The partial physical instantiation concept requires a new type of interface. We conceived a generalized object outline, similar to a string. The outline of the string would correspond to the profile of the virtual object at a given 2D slice.

While the string interface shuns the use of technologically specialized tools, any physical object can be applied to the string as a tool. Profiles can be introduced to the string with CD cases, pens, cell phones, childrens blocks, or cookie cutters. Patten classifies this level of improvisational interaction "object as form", using the physical shape of the items [32].

Beyond the improvisational tools, the SP3X loop was designed to engage the most commonly-used and important tool the hands. The low layer of division in SP3X results in high bandwidth; both hands and all ten fingers can be engaged simultaneously. The ability to use both hands and many fingers is expected to lead to many new interaction styles.

A device for freeform creation requires positional freedom. In order to maximize this aspect, a structure with the full six degrees of freedom is desirable. The decision to use a structure at all was heavily influenced by Fiorentino [9] and the identification of both fatigue and precision as central problems in existing VRAD solutions.

The structure provides a spatially-positionable working surface. The platform needs to be supported by the user only when it is being moved to a new location. Most of the time, the user only needs to support his or her own arms, and does not have to hold them in a particular position or orientation. In addition, the user is able to rest the wrists on the platform, permitting greater precision in object manipulation.

The design of SP3X separates gross and precise activities. The interface encourages and even requires the user to use the physical structures that are best suited to the precision of the task. Gross navigational manipulations are handled by the shoulders and arms when the platform is flown to a new location. The arms and shoulders have a broad range of motion and strength, but relatively coarse control. This allows the user to easily move the platform anywhere within their reach. Precise sculpting of the contours is left to the hands, which shape the curvature of the loop on top of the platform.

The moveable working surface maintains spatial coincidence between the loop and the virtual model. The alternative, non-coincident, could be executed by building a fixed working surface like a table and a location placer tool. Coincident interfaces are more intuitive because they visually demonstrate the association of the two elements.

4.2 Display

We considered several display options. The governing principles were support of:

visual navigation: ability to change the viewing perspective easily

continuous visibility: output model should always be visible

spatial perception: conveying the location and shape of the object

Each of these contributes to the overall usability of the system. Visual navigation eliminates the need for additional navigation tools or modes. Continuous visibility encourages an experimental trial-and-error approach because the user receives continuous feedback. Spatial perception is essential because the goal of the device is the production of an object whose primary impact is visual.

In order to permit visual navigation while maintaining the basic coincident principles of tangible media, the display itself must be moveable. Previous research systems like the Boom Chameleon [41] and the magic lens of Metadesk [43], as well as the commercial Magic Window [SquidLabs], make use of a tracked-location display that is manually positioned. This requires the user to move both the working platform and the viewing window, and it would be very difficult to move them simultaneously. The broken navigation/interaction configuration directly interferes with continuous visibility and hinders trial-and-error experimentation.

Canon mixed-reality HMD.

We finally settled on the use of a head-mounted display. This system provides us with invisible navigation through head tracking, spatial perception through stereo vision, and continuous visibility by mediating anything that the user can see. The Canon HMD which was chosen is also designed for mixed-reality, allowing the user to see the physical device directly while superimposing the virtual object.

4.3 Early Prototypes

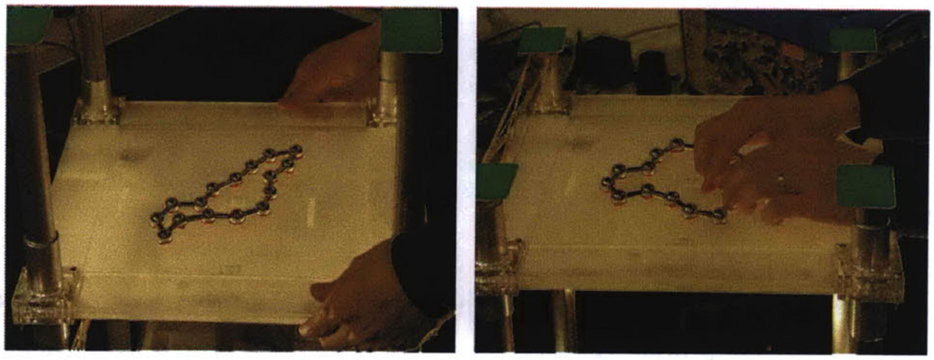

Figure 4-1: The Manipulable Irregular Solid Creator

We created a prototype device, named the Manipulable Irregular Solid Creator (MISC), to explore the potential of partial physical instantiation and planar definition. MISC, as its acronym suggests, was conceived with the lofty goal of taking the burden of creating miscellaneous objects from the shoulders of GUI-based CAD. The MISC system had a limited working space, constrained to a vertical column. A square working surface translated along four vertical pipes to yield a rectangular prism volumetric working space.

The interaction was straightforward- the user sculpts a chain of ball bearings and rod magnets into a closed loop of some shape, and that outline defines a two-dimensional plane within the three-dimensional working space. The working surface is then translated to another vertical position, and another plane defined. Interpolation between these two consecutive planes results in a 3D solid.

After informal experimentation, we decided that the MISC systems constraints encumbered the user too greatly. The rectilinear nature of the structure made it impossible to create models of the free-form type we had envisioned. Users commented that they did not have the freedom necessary to do what they wanted, and would try to rotate the platform.

We investigated structures with more degrees of freedom, including arms and Stewart platforms. After deciding on a hexapod, I constructed a LEGO model of the structure to explore the physical and mechanical limitations of the device and to visually determine the effect of differences in the ratio of deck to base dimensions.

4.4 Implementation

4.4.1 Mechanical Design

The mechanical characteristics of the system would determine both the overall feel of the interaction and what capabilities were possible. I decided that the physical structure needed to meet the following criteria:

Full six degrees of freedom: maximize 3D freedom of motion

Compact: small lateral footprint to avoid obstructing the user

Smooth motion: fluid motion to encourage free-form results

Wide range of motion: operational area exceeding human reach from standing position

Force feedback capable: can be augmented with actuation to support physical feedback

Easy, accurate instrumentation: aids measurement, rather than making it difficult

Cost and Manufacturability: relatively easy and inexpensive to build

After a great deal of deliberation and research, a Gough platform was chosen. This six-legged structure commonly actuates motion simulators and small 6 DOF milling machines. A Gough configuration (or hexapod) brings several advantages, particularly in comparison to multi-segment arms.

The hexapod segments work in parallel, rather than in series like the arm. This structural difference means that each leg-joint unit only needs to support a share of the platform. In the arm configuration, each joint must support the load of all the more distal segments. This mounting weight, coupled with the great leverage of a long arm, forces them to be designed with segments decreasing in size moving outward, and in the case of force feedback, motors increasing in strength moving inward toward the connection point. The long lever arm becomes a great concern when considering that a human user supplies the end-of-arm load.

Arms also suffer sensing accuracy problems. The end position is calculated serially from rotation and angular sensors at each joint. Therefore, any error in the joint position nearest the connection will result in an error in the assumed location of the second joint. Minor errors at each joint compound and lead to serious inaccuracy in the computed location of the final segment. From a design perspective, these concerns necessitate the highest accuracy sensors available and inflate the cost accordingly. In contrast, the hexapods struts are instrumented and operate in parallel, so measurements are completely independent.

The hexapod brings a final advantage: redundancy. Each of the six legs is identical, and this saves time both in design and in manufacturing. It also improves ease of repair, as parts can be substituted, and complete replacement units swap for any damaged component.

Hexapod Joints

The hexapod contains a total of twelve rotational joints, two per strut. Every joint requires at least two degrees of rotational freedom, which makes ball joints and universal joints suitable. Ball joints of sufficient size carry load well by transferring the force directly into a socket. The arc of the shaft opening and the diameter of the shaft determine the range of motion. The "smoothness" of the joint depends on the compatibility of the materials, contact area, and lubrication. Universal joints generally find use in applications with angularly offset rotating shafts, like truck axles. U-joints carry force through pins, and are not designed to support load.

U-joints are available in lightweight plastic form, while ball joints are most commonly steel for wear resistance.

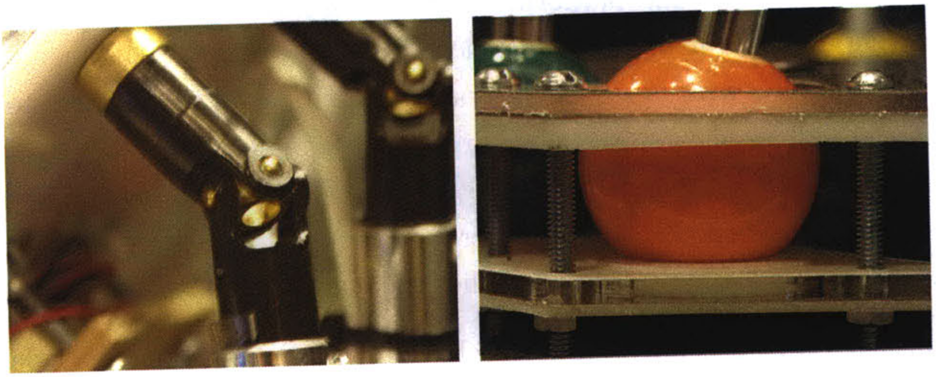

Figure 4-2: Universal joint (left) and ball joint (right).

I chose to use both types of joints. Large ball joints in the base carry load smoothly and easily with a massive range of motion, while lightweight u-joints contribute low load and free motion at the deck. The ball joints, if designed to provide side access, could be augmented with brakes for passive force feedback.

Extensive investigation proved fruitless for large ball joints with a wide swivel range. McMaster-Carr, for instance, stocks ball joints with a maximum angle of less than 40 degrees. By my calculations, to encompass the width of easy human arm extension at waist height, a minimum of 90 degrees is required. Also, commercially available ball joints completely enclose the ball, eliminating the possibility of joint braking.

Therefore, I designed and manufactured my own. I developed a 2.25" diameter ball joint based on a billiard ball with a PTFE (Teflon) seat and capture. The seat and capture were waterjetted out of 0.25" PTFE plate and then rounded with a blade. One comment on waterjetting Teflon- the material is hydrophobic, and forms a sheet of scum on the surface of the water which is bothersome to skim. In order to save on expensive material, the Teflon components were concentric, so the seat and capture fit in the same area, with less wasted material. The seats pressure-fit into acrylic sockets, further reducing the Teflon needed. The assembly securely holds the billiard ball and allows 120 degrees of motion with the chosen shaft. Bolts compress the joint, and allow the resistance to be altered for different holding force. This changes the "feel" of the platform.

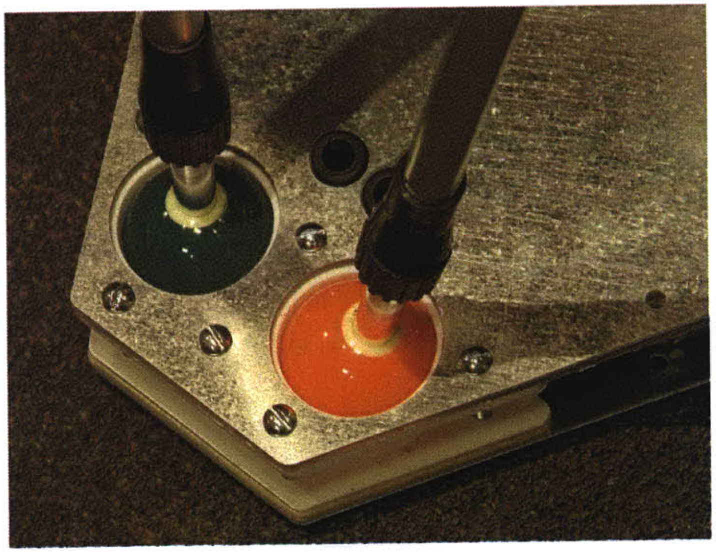

The billiard balls were drilled to snugly fit a steel tube. Centering the bore proved challenging, and any offset would impair the function of the joint. To ensure accuracy, a board was clamped in place on the drill press, and a hole saw used to remove a perfectly centered seat. Without unclamping the board, the hole saw was exchanged for the correct gauge bit, and a billiard ball was secured in the seat. This method yielded excellent results.

The universal joints were initially selected for weight and price. Early experiments to manufacture them went poorly, so commercial sources were sought. Nylon joints met the criteria- both inexpensive and very light. As an added advantage, nylon mills easily. These joints functioned acceptably, but occasionally separated- to the surprise and annoyance of the user. Eventually, under wear conditions for which they were not designed, the joints began to break. After replacing one, an alternative was needed. A much more expensive type composed of Delrin and brass was found. These have no separation problems and are much smoother. They are slightly heavier, but more durable.

Hexapod Struts

Six telescoping segments connect the ball joints in the base to the universal joints in the deck. These must slide as smoothly as possible, and resist off-axis forces with minimal frictional increase and zero jamming. Telescoping rods from industrial suppliers were both heavy and expensive, so we sought consumer items that could be repurposed. Camera tripods, extending mops and dusters, hose attachments, and more were evaluated for the desired features. Potential jamming problem was the primary reason for elimination. A garden hose attachment designed for cleaning gutters proved impossible to jam, and possessed both the smoothest motion (due to O-ring bearings) and a rotating cuff for locking the position. The latter feature would become important in the design of the brakes.

Hexapod Brakes

Early in the design discussions it became clear that force feedback held powerful potential for interactivity in the support of navigation and simulations. Force feedback (FF) could be executed in a variety of ways on the hexapod platform.

Active vs. Passive

Active force feedback involves an actuator that directly exerts force (and imparts motion) on the user. In order to be effective, an active FF must apply forces that are similar in magnitude to the users strength. This introduces significant energy requirements, and more importantly, creates a potential safety hazard. A malfunction in the control system of a robot with human-scale strength could easily turn user into victim. Active force feedback is widely in use on very large hexapod systems used for aircraft simulators.

Passive force feedback is a system through which the illusion of force is manifested by selectively retarding the motion of an object. In contrast to active force feedback which creates motion, passive FF creates “stopping”: it is a braking function. The actuator in passive FF works on an object, rather than on the user, so there is no direct path of energy to the user. This structure is inherently safer than active FF.

Brake location

The hexapod has six degrees of freedom, and all six of these must be controlled for effective feedback. With passive as the chosen method of force feedback, the required brakes could be mounted either on the legs or around the joints.

Strap-type brakes within the base of the hexapod could be used to limit or eliminate rotational freedoms of the struts. The base would completely enclose these brakes, making for a cleaner appearance. However, they would require enough frictional force from contact area and actuator power to overcome human strength exerted at the end of a 1+ meter lever arm. This requirement is very difficult to meet with actuators compact enough to fit within the base. A more destructive problem is that each ball joint has two degrees of freedom- so braking a single joint results in the loss of two degrees. This would make it impossible to control the platform completely, as any FF action impacts two degrees simultaneously and equally.

Compressive brakes on the individual struts present a more compelling solution than strap brakes. Each acts only on a single telescoping segment, and thus controls a single degree of freedom. Another benefit is that the lever arm problem does not occur, as they are positioned on a segment rather than on a joint, and this reduces the force required.

The greatest advantage to this approach is the fortuitous construction of the telescoping struts. In a holdover from their former lives, these struts each come equipped with a rotating constriction brake. This eliminates the need to fabricate such mechanisms from scratch. Moreover, the rotating collar gives range of braking resistance over a narrow range of motion. The short actuation distance promotes fast response times to any of the resistance levels, from completely free to fully stopped.

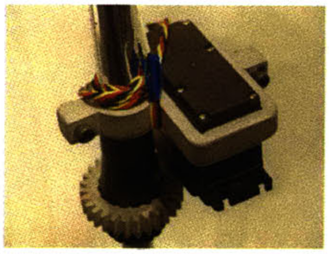
Servo actuated brake.

Hexapod Sensing

Measurement of the lengths of the six struts and subsequent calculation permits the accurate determination of the location of the deck. The quality of the length measurement dictates the accuracy of the deck calculation, so this sensing is critical.

The struts travel approximately 32 inches, at maximum. Any sensing technique chosen for the struts must therefore be capable of covering this distance with high precision and accuracy. Retracting string-based linear encoders solved these problems but, at $200 each, were relatively expensive in a scenario requiring six.

We chose to solve two problems with one device. The system required both length measurements and electrical communication to servos residing on moving segments. By measuring changes in the length of a multi-stranded wire attached to the servo, the wire can serve both functions.

We chose a retracting telephone cable unit to serve as the core of this structure. The retraction feature operated as a built-in cable tensioner, and the RJ-45 connectors simplified integration. The cable unit was heavily modified to incorporate new mounting hardware, wire angle rerouting, and custom laser-cut optical-encoder wheels. These encoder wheels were designed to match the spacing of USB mechanical mice. Remounting and wiring the elements from the mice equipped us with a nearly ready-made sensor package.

4.4.2 Electronics

The brakes are powered by large servos (CS-80) which require basic control circuitry. The CS-80 is a Giant Scale servo that can produce 20+ kg-cm of force and consume 2 amps or more at 5 volts. Servos have the advantage of separating the lines for power and control. This simplifies the control electronics dramatically, because we can connect a microcontroller directly to the servo without intermediate MOSFETs to handle the high currents.

Another advantage to servos is that they have built-in position sensing. This feature improves repeatability over systems without feedback loops, and greatly simplifies the construction over systems with feedback loops, which do require sensors. Servos use pulse-position modulation (PPM), a control communication standard in which the length of pulses correspond to the desired position of the servo. The range of pulses is 1-2 milliseconds, with a pulse of 1.5ms corresponding to the center of the servos range of motion.

The PPM signals are generated by an Atmel AVR ATMega32. Each of the

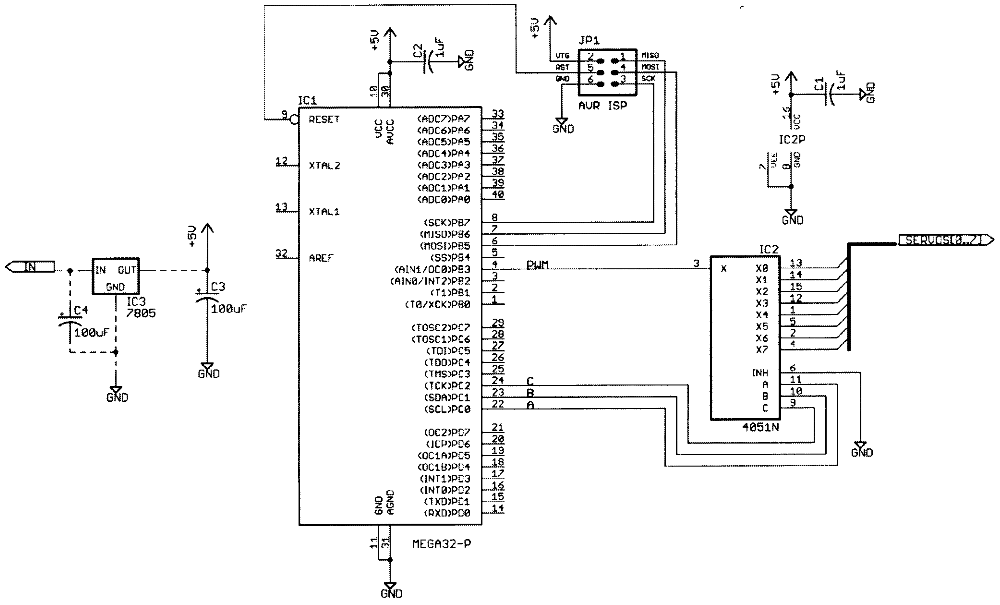

Figure 4-3: Circuit diagram of the control electronics for the servos.

six servos can be addressed individually with a theoretical resolution of 32 bits. In the current configuration, the AVR detects the state of a lever switch mounted to one of the platform handles. If the lever is depressed, the servos simultaneously move towards the end of their movement range that "opens" the brakes. When the lever is released, the servos return to the default "closed" position, relocking the brakes.

The power for the servos is supplied by a standard 300 watt AT power supply, originally designed for computer usage.

4.4.3 Software

The software behind SP3X can be divided into three categories. The first is the computer vision and tracking software that makes the camera capture of the loop possible. The second is the application that captures input from the length encoders and processes it into usable information. The third and final segment of the software is the geometry and rendering engine that combines the data. My esteemed collaborator, Mahoro Anabuki, wrote both the vision tracking software and the geometry and rendering engine.

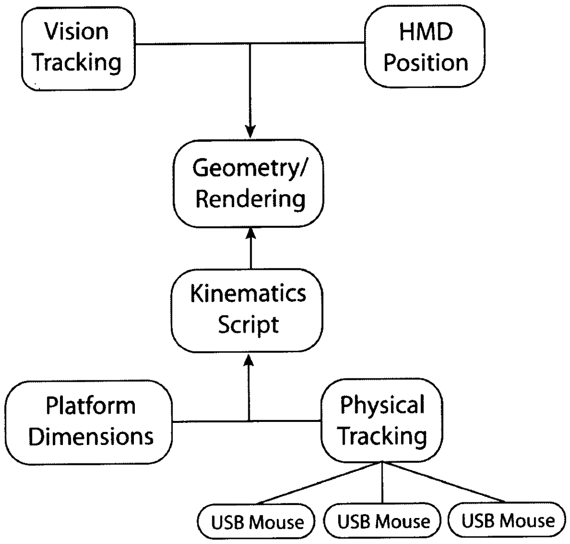

Figure 4-4: Simplified model of the software architecture of SP3X.

Vision Tracking

Computer vision is a method of automatically identifying features in an image using processing algorithms. In this case, the image is captured by one of the two cameras in the HMD. The image is then analyzed looking for specific tags. The tag we are using is a shade of the color green. Pieces of green paper have been taped around each segment of the sculpting loop, and these tags make the segments computationally simple to identify. The algorithm ignores the red and blue data in the image, and then uses a threshold to eliminate background noise. The threshold works by excluding any data that is darker than a certain threshold value. Then, groups of green pixels are identified. The centers of the green blobs are calculated using a center-of-mass algorithm, and the coordinates of the centers are recorded.

Physical Tracking

We determine the location and orientation of the platform by measuring the lengths of all six legs and then applying forward kinematics to the data. The leg extensions are based on data from the modified mouse encoders. Computer mice output values based on the distance traveled during the last measurement interval. Since this is not an absolute location, the positive and negative distance outputs must be stored in a variable and a running sum maintained. We convert this sum of mouse "tics", which have no distance-unit equivalent, into centimeters through a multiplicative factor.

The final resulting six numbers correspond to the lengths of each hexapod leg. These lengths are used as input to a forward kinematics script written by Merlet [28]. This software returns the location and orientation of the platform in relation to the base.

Geometry and Rendering

The final block of software brings together the data from the previous two. The centers of the green blobs are chained together to form polygons. The order of the vertices is determined by selecting a vertex, comparing distances to other vertices to find the closest one, and then continuing this process around the loop. The planar polygons are then oriented in the three-dimensional space by applying the platform location using a transformation matrix. Sequential planes are connected vertex-to-vertex in numeric order to produce solids. The resulting data constructs are rendered using OpenGL.

4.5 Concerns and Limitations of Implementation

The design and implementation of the SP3X system was the product of much collaboration and compromise. Some of the characteristics of the resulting device raise concerns regarding both the output and the interface. Three of the central concerns that have arisen are that the output will be restricted and simplified by the use of coplanar points, that the separation of the 2D and 3D modes will impose a time cost in transitions, and that the loss of input/output coincidence between loop instances will be problematic.

Simplicity of Output. The structure of the SP3X device, in which a flat sheet of steel is used as the work surface, constricts the points of each cross section to a plane. The necessity of coplanar points definitely reduces the complexity of the objects to some degree. The proportions and resolution of SP3X (work area to working space) also encourage columnar forms. In general, SP3X most readily produces objects that are axially simple but peripherally complex.

Separation of Modes. A typical usage scenario involves both 2D manipulation on the work surface, and the 3D navigation of the work surface through space. These interactions cannot reasonably be performed in tandem, and the need to transition from one hand position to another imposes a time penalty. Informal observation has lead us to believe that the time loss is 2-3 seconds, timed from release of the loop to first motion of the platform. One second of this delay can be attributed to the release time of the brakes, as their range of motion has not been optimized, and users wait until the motor sound has ended to begin moving the platform.

Loss of Coincidence. It is impossible to revisit a captured slice. When a slice is created and the platform navigated away, the physical loop loses coincidence with the virtual slice. Returning the platform to the same position and orientation is possible, but the shape of the loop will no longer match its virtual counterpart. Successive slices are independent instances of the 2D interface, and the instance is broken upon capture. It is possible to solve this issue by using an actuated loop, but the design of a feasible solution has eluded us.

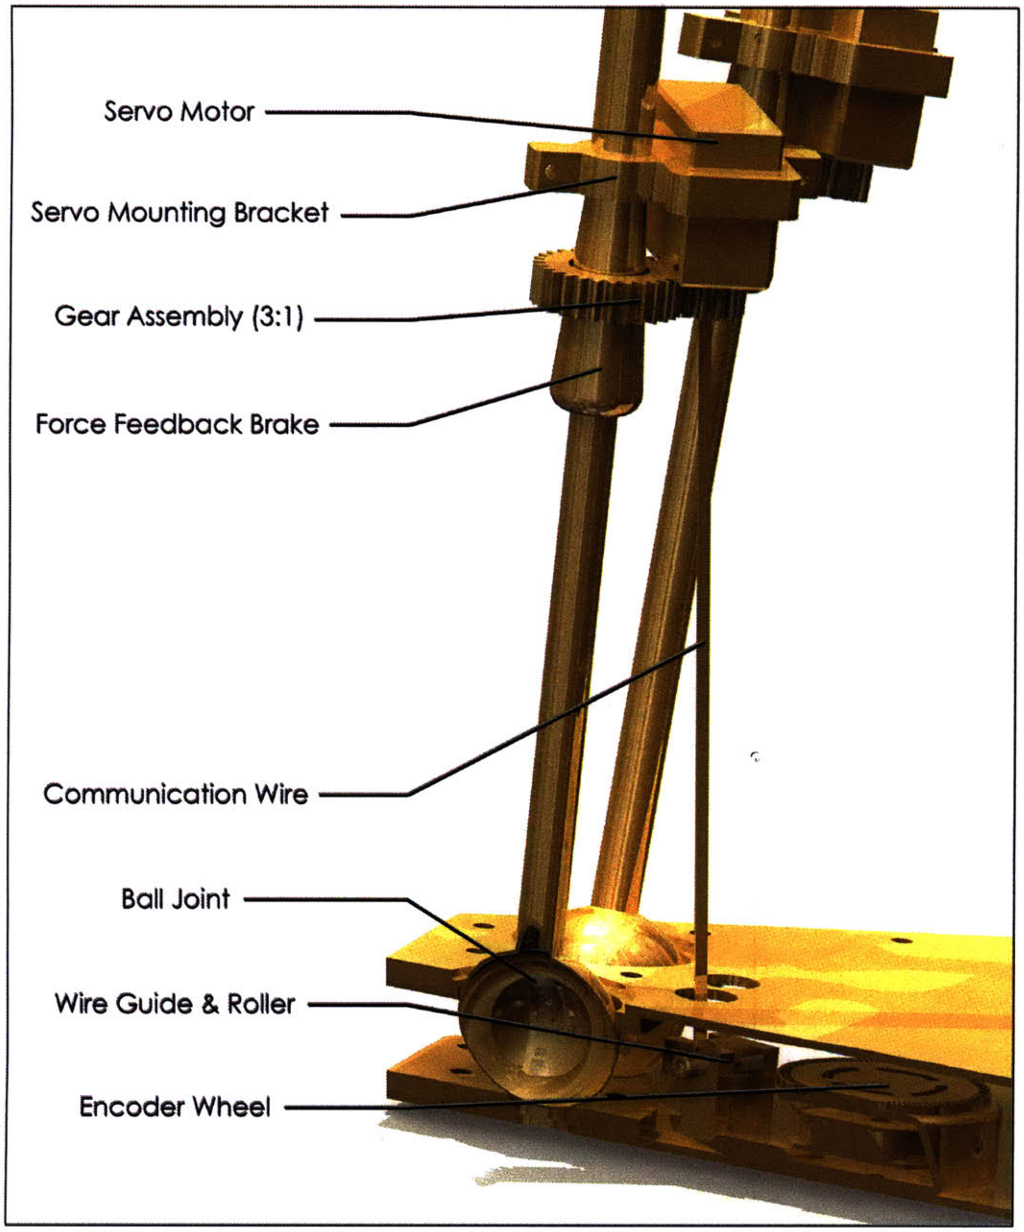

Figure 4-5: Render of the final structure and components.

Chapter 5

Evaluation

5.1 Introduction

Experiments were designed to observe the differences in usage between the new device and GUI-based CAD. The experiments described in this section were structured to uncover whether users took advantage of the capabilities of SP3X, and what interaction languages they created.

In this experiment, participants were asked to perform three tasks. The first task was designed to ease the participant into using SP3X and to observe the learning process. The second task was the sculpting of an abstract idea, to delve into the artistic potential of the interface. The final task served to evaluate the participants ability to use SP3X to accomplish a specific task: recreating a given structure. These tasks were designed to explore SP3X as an easily learned modeling tool.

5.1.1 Apparatus

The apparatus used for this study was the SP3X platform described in chapters 2 and 4. Each participant was asked to use the hexapod device in three modeling tasks while they were observed and their digital output recorded.

5.1.2 Participants

A preliminary study of the interface was conducted in August of 2006. The participants were MIT undergraduate and graduate students. A total of 10 students responded to a general email to several groups. The group was comprised of 7 males and 3 females, spanning 19-33 years of age. While it was not requested, 80

5.1.3 Experimental Procedure

The experiments use three scenarios: introduction and simple tasks, artistic creation, and sculpting a letter. The introduction scenario asked users to follow written instructions to acquaint them with system basics. During this scenario they were free to ask the proctor if they became confused. The artistic creation scenario encouraged the user to explore abstract forms and create a digital model of it using SP3X. In the letter sculpting scenario, participants were asked to create a model based on a single letter, possibly one of their initials. All users participated in all three scenarios.

The first task, the *introduction*, was created to observe how quickly participants became familiar and comfortable with the interface. For the first 10 minutes, the users were given the HMD and allowed to experiment with the device without instruction. The users were then given a set of written instructions that walked them through very basic tasks like moving the platform and creating an object. They were asked to complete these instructions as quickly as they felt comfortable. After completing the instructions, participants were allowed to experiment for another 5 minutes if they wished.

The second task, *artistic creation*, was designed to document how effectively participants could apply their new knowledge. Participants brainstormed a form, and then were asked to use SP3X to make a digital model of it using a maximum of 10 minutes. This scenario was composed to explore whether SP3X was a useful tool for ideation and on-the-fly modeling.

The third task, *letter sculpting*, was used to test the potential use of SP3X for creating recognizable forms. In this task, the participant was asked to create a three-dimensional representation of one of the letters in his/her initials.

5.2 Results

We received very positive feedback overall, with users quickly adapting to the interface. One participant felt that there was perhaps too much freedom, and "found its flexibility a bit daunting." The participants were impressed by the potential of the interface, and many suggested applications and offered ideas for improvements.

One surprise from the introductory task is that most forms were constructed of parallel planes and were relatively columnar. We believe that is most likely the result of early stage exploration of the system, and ties into the "daunting" nature of the freedom. This behavior continued into the artistic task; though users were more likely to create non-columnar forms, the planes were commonly parallel. When encouraged to deviate from the safest and easiest forms by the letter sculpting task, subjects produced more planar rotations. The escalating form complexity supports the belief that early column structures are a form of exploration by reducing the number of variables.

One participant noted that the basic usage of SP3X was obvious and easy, but that it was also similar to learning any new tool and would require some practice to become skilled. He also called the system "forgiving," and said that the preview helped plan objects.

A particularly observant participant commented on the differing levels of precision in platform movement and loop sculpting, and quipped about the discrepancy.

5.2.1 Components

HMD

The head-mounted display drew the most flack from participants. It was criticized for being uncomfortable and encumbering, while the cable was seen as a distraction and was "perpetually in the way." Wearing the goggles was also listed as an obstruction to collaboration. One of the users put his problem very bluntly: "I cant imagine wanting to put the silly thing on."

What was more surprising is that two participants suggested that a 2D display on a monitor, seated to the side of the whole affair, would be more than adequate. One participant, with prior experience in sculpture, spoke of his ability to visualize the object, without the need for a display device. He stated that he could visualize where he had created planes accurately, and found that the HMD simply got in the way.

Brakes/Support

The locking feature of the platform was broadly praised, and even cited as "key" and "awesome." The users loved that there was a stable working surface. A participant noted that the brakes provided auditory feedback, and that this was a very useful feature. Despite these benefits, it was commonly complained that the structure was too heavy and was cumbersome to lift.

Loop

Participants made many positive statements when asked about the interface loop. They found that the "bead string was very intuitive" and that it was "easy to understand," both of which support our goal of an intuitive system.

Granularity

One participant complained about the low resolution of the loop, and inquired as to whether it was possible to double the number of segments to create a smoother curve. When assured that it was possible, the participant went on to say that while the potential of the loop interface was clear, the low segment count of the loop made the output feel clumsy.

Handles

We found that participants almost always gripped the handles from the top in an over-handed fashion, rather than from below in a lifting motion as we had designed. This made the platform clumsy to translate vertically. We believe this problem is probably due to the initial height of the platform. The top surface was slightly below elbow height for most subjects, so they reached down in a straight line, rather than reaching underneath to lift. This might be improved by raising the resting height of the platform or by redesigning the handles so that they are more easily grasped from below, though the latter solution would certain become an annoyance in certain orientations.

The placement of the brake release button proved awkward, as it was designed for use with an underhand grip. Participants suggested that a button placed in middle of the handle would be more accessible. The original design had been exactly this configuration, but it was discarded due to concerns about accidental activation.

Pedal

We observed that the current pedal, used for triggering slice capture, had several problems. As an industrial machine shop pedal, it has a foot guard that prevents accidental activation. On their first attempt, most users tried to depress the guard instead of the pedal. Users did not receive clear tactile feedback from the pedal, and appeared to rely heavily on the display to confirm that the slice had been captured. By far the greatest issue was that the pedal anchored participants. They retained contact with the pedal at all times, greatly reducing their movement freedom.

5.2.2 Observed Interactions

During the course of the experiments, we observed several types of loop interaction:

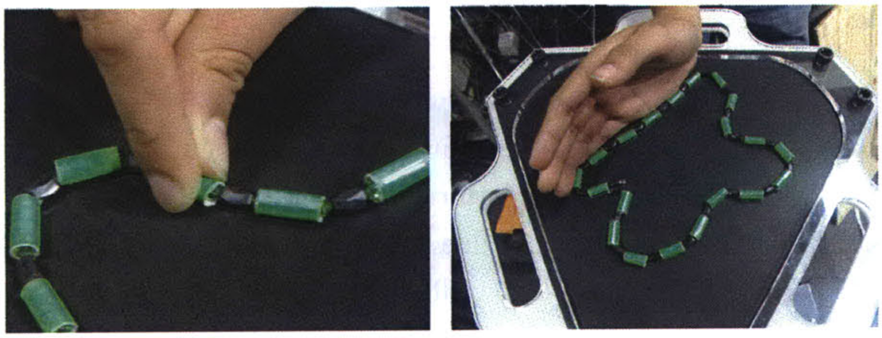

Figure 5-1: Pinching (precision modification) and Bulldozing (group modification)

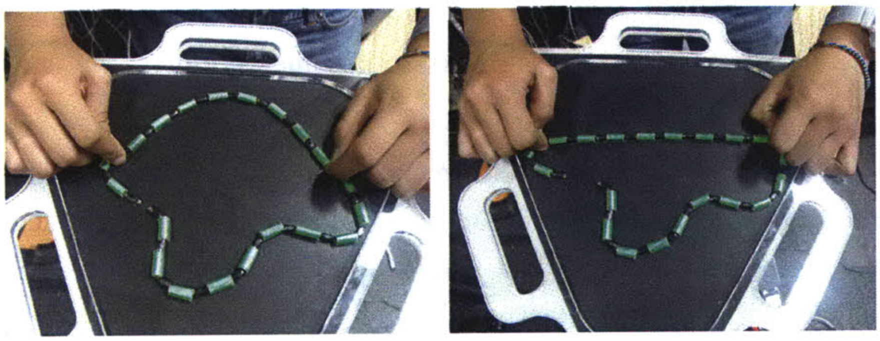

Figure 5-2: Stretching Creating straight lines using physical limitations

Figure 5-3: Imprinting Using the hand to create and transfer profiles

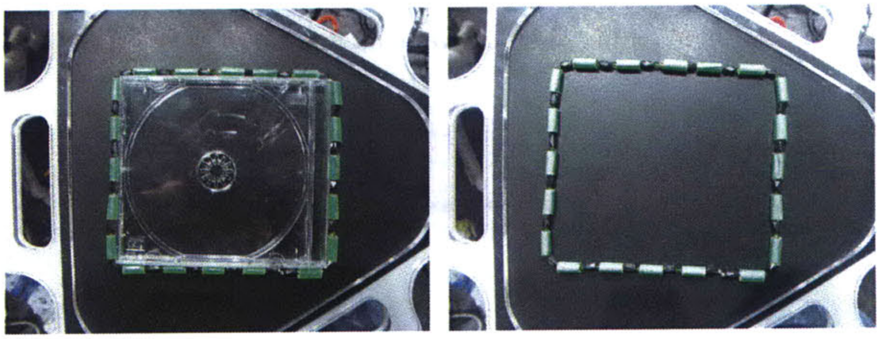

Figure 5-4: Matching Using the shapes of objects to guide profiles

5.2.3 Artistic Creation Results

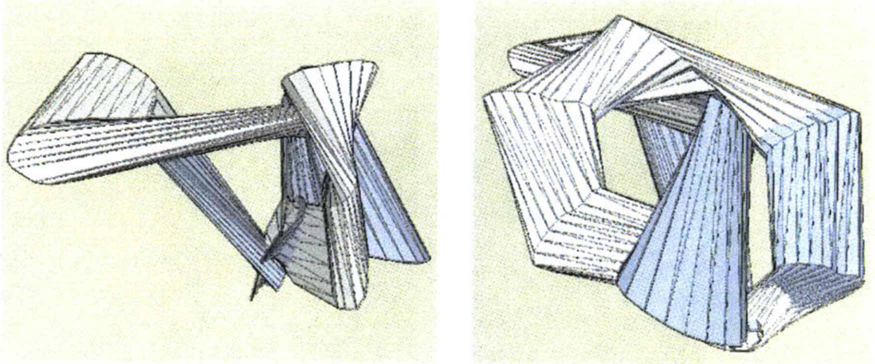

Figure 5-5: "The Knot", as rendered in Google SketchUp.

Participants generally spent more of their time on the careful creation of cross sections during this section of the study. The two exceptions, who focused more time on the overall structure, produced what we believe are the most visually interesting of this section. The knot and the double helix appear complex when rendered from many angles, while the Pacman is relatively two-dimensional.

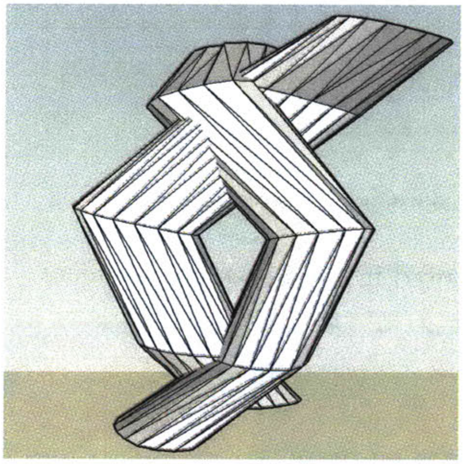

Figure 5-6: "Helix", as rendered in Google SketchUp.

5.2.4 Letter Sculpting Results

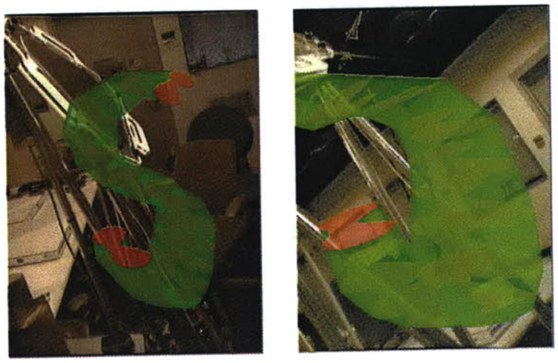

Figure 5-7: In SP3X capture of the letter "S".

The same letter "S", exported to Google SketchUp. While the letter is clear from the front, the side perspective shows irregularity.

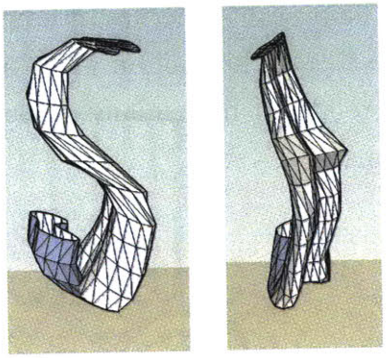

Figure 5-8: Google SketchUp render of the letter "S".

The letter "J", captured in the SP3X environment.

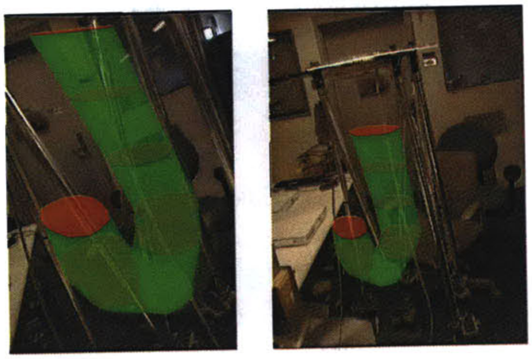

Figure 5-9: In SP3X capture of the letter "J".

With a single exception, the participants went for the large scale in the crafting of their letters. This focus on the 3D navigation element of SP3X may stem from growing confidence with the interface, impatience to finish the study, or perhaps a desire to draw a letter from their name very large.

5.2.5 General Suggestions

The participants proposed that SP3X might be well-suited to some applications. Participants mentioned architecture, landscaping, and animation, fields in which we had already identified potential. One subject had been overheard during the study saying, "This thing could be great for designing climbing walls," and later suggested that same application, citing the irregular surfaces that were popular.

One participant stated that it would be preferable to have a wall-mounted unit, with usage similar to arm-mounted desk lamps. SP3X could be stored in an "up" position, out of the way, but still easily within arms reach.

Another user observed that the display of virtual object was too transparent, and that this led to confusion. It was suggested that the object be rendered completely opaque with studio lighting. When asked about occlusion, the user scoffed and said that he would simply look around the other side, as though this were obvious and natural.

5.3 Summary

From the experiments it is clear that the SP3X system possesses the potential to be a useful tool for beginning modelers. First-time users, both those with and without prior CAD experience, found the interaction easy to learn and were able to produce forms that, while primitive, are recognizable.

Investigating the ongoing usefulness to professionals and serious amateurs will require a long-term study involving regular work with the device.

Many alterations and tweaks would be necessary to move SP3X out of the research environment and make it commercially viable. These will be covered in the Future Work section of the next chapter.

Chapter 6

Conclusion and Future Work

6.1 Conclusions

The research presented in this thesis evaluates the application of a novel device using six degrees of freedom and implementing the philosophy of partial physical instantiation for three-dimensional modeling. The strength of the SP3X system lies in leveraging the strengths and weaknesses of human perception and physiology. SP3X draws on the research of both mixed reality and tangible interfaces to yield a new solution to the digital clay problem. Experiments proved that the resulting system is intuitive, supports a varied range of interactions, and outputs recognizable and usable digital models.

The main contributions of this thesis fall into two categories. The first and primary contribution is the extension of hybrid interfaces into the partial physical instantiation of objects, the result of merging elements of mixed reality and tangible interfaces to produce a new flexible and inexpensive approach to three-dimensional display and interaction. The secondary contribution is the interface itself, consisting of a sculpting loop and hexapod-supported working surface. The combination of these contributions results in an interface with a demonstrably low floor and the strong potential for a high ceiling.

The scope of our research encompasses determining whether a specific style of interface, incorporating the sculpting loop and partial physical instantiation, can be intuitive and powerful for novice users.

The SP3X device was designed not a substitute for traditional GUI-CAD, but as a peripheral augmentation and as an option to novices for the creation of complex solids. In comparison to GUI-CAD, SP3X lacks almost all of the functions that make CAD indispensable to engineers and professional product designers. In place of these mathematical functions, SP3X offers a simple interface for creating geometrically irregular objects with a minimum of obfuscation. Therefore in its current iteration, we recommend that use of SP3X be limited to artistic expression, unless the output is parametrically refined in traditional CAD.

Low Floor User studies provided insights into the human responses and usage patterns of an unfamiliar device with a high degree of dimensional freedom. The interface was shown to be intuitive and quick to learn for first-time users by supporting natural behaviors and physical expectations. These behaviors are a result of a lifetime of experience with basic physical objects, and SP3X attempts to meet these expectations.

Human subjects who were new to the use of this device demonstrated confidence using this system within a few minutes of exploratory use, and were able to create surprisingly complex objects. They also spontaneously developed a variety of interaction techniques with the sculpting loop, without encouragement or prompting.

High Ceiling In the future, a longer-term study would be required to ascertain the "high ceiling" potential of this interface. We believe that SP3X has great potential power due to the human emphasis of the design. Much like artisan tools such as a potters wheel and a sculptors chisel, the quality of the output is dependent on the skill of the artist. More importantly, this ability is based on the cognitive skills of imagination and planning and on the physical skills of dexterity and coordination. Just as these may be cultivated with practice and experience, we posit that artists can improve in the use of the device.

6.2 Future Work

The design decisions of SP3X were selected with the goal of producing a working system that embodied the central philosophies of the researchers. Contributing goals included minimizing cost and the number of custom parts, and reaching a working prototype quickly. In doing so, many compromises and choices were made which greatly affected the resulting device. Not all of these solutions were ideal, and here we present improvements and alternatives to the current configuration.

6.2.1 Improvements

The following future changes are suggested if further development of this prototype is pursued:

Visual Tracking. The present implementation of the visual capture of the loop has distinct drawbacks. The camera used for capture is located in the HMD, and this introduces issues with occlusion, complicates the computation necessary to locate the loop in 3D space, and creates variations in both lighting and resolution of the loop. We propose that this could be solved by mounting a small camera on the edge of the platform. This will insure both an unobstructed view and a relationally static viewpoint. The addition of a dedicated light source, coincident with the camera, will reduce issues with orientation-affected ambient lighting. Another improvement under this configuration would be the use of retro-reflective tape as markers on the loop, increasing their illumination over the ambient.

Reduce Platform Weight. The platform, while readily moveable, weighs enough to be clumsy at certain angles. The majority (70

Granularity. The current configuration captures a maximum of 20 points per plane. While this resolution is sufficient for early investigations, it limits the users and may also encourage them to produce overall simpler models. This can be fixed by modifying the software to search for more points, creating a loop with a higher density of segments, and insuring that the camera will have adequate resolution to pick out the smaller segments.

Force Feedback. All of the necessary elements for simulation-linked force feedback are present in the system. The implementation of an application which requires it and the force feedback software remain to be created.

Displays. Subject feedback on the HMD has been overall negative. Investigating alternative displays may prove to be a fruitful endeavor, and we recommend trying displays of the desktop, arm-mounted, and coincident varieties.

Actuated Loop. The loss of loop/model coincidence could be solved by the implementation of a mechanically actuated loop that maintains the correct profile as the platform traverses the virtual model. This option was discussed at length during the design phase, but all of the suggested solutions imposed penalties that impaired the basic usability of the system.

6.2.2 Alternative Structures

The human-scale hexapod platform chosen for several reasons (discussed in Chapter 4), however other designs may be useful to others:

Wall Mounted. One of the study participants suggested that a device of similar proportions could be mounted on a wall to free up floor space. In his design, the structure would fold flat against the wall when not in use. When needed, it would extend out into a work volume. Though the position of the wall limits the users range of perspectives, our study indicates that people tend to use only 180 degrees currently, so this does not present a true impediment.

Smaller Scale. A desktop version of the device seems more appropriate for personal applications. The miniaturized SP3X would occupy the space of a mousepad, and thus be more convenient for people without the floor space to dedicate to the device. This smaller scale will also reduce the overall cost and make it more commercializable.

6.2.3 Applications

The modeling application presented in this thesis demonstrates the first glimpses of the capabilities of the SP3X system. The following are other applications which have been suggested.

Oil Well Pathing. Oil wells can be modeled as long splines, often with turns and twists. These splines follow a complicated path through the diverse strata of geologic formations. Using present technology, this course is impossible to calculate automatically, owing to the myriad of contributing risk and cost factors. The wells are laid out by hand in an unsupported virtual environment. The SP3X system could be used not only to reduce fatigue, but also to supply force feedback cues based upon the risk and cost factors.

Product Housing Design. The force feedback capabilities of SP3X could also be applied to the creation of shells for products. While this is largely an aesthetic endeavor, product housings may be sculpted around existing components. In this case, force feedback can alert the user to overlaps and intersections with vital internal elements. The already shell-like form of the loop further supports SP3Xs potential use in this field.

UAV Control. The six degrees of freedom that SP3X offers matches well with the range of control necessary to manually fly a robotic helicopter. Unmanned aerial vehicles (UAVs) of the aircraft type are also suitable, though they require fewer degrees of control to fly.

Video Games. In a very similar vein to UAV control, video games are a very promising application area. Spaceflight games, in particular, would be well-suited to SP3X. Instead of separating out the six degrees over a mouse/joystick and keyboard, they could be concentrated into a single device whose orientation is directly mapped to that of the vehicle. Two desktop SP3Xs could even be used in tandem to fly two craft.

Music. Two members of the Media Lab, upon playing with SP3X for the first time, commented that it could easily be used for electronic music generation. The ability to control six variables independently and fluidly is appealing, as is the performance aspect of playing complicated electronic music through a single device using a gestural dance.

6.3 Closing

We hope that the work presented in this thesis may support the efforts of contemporary and future researchers who also pursue the goals of tangible interfaces, easier computer-aided design, and virtual clay. Promoting these research areas may eventually lead to the user-designed mass-customization future which we envision.

Bibliography

[1] Raffaele De Amicis, Giuseppe Conti, and Michele Fiorentino. Tangible interfaces in virtual environments for industrial design. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, pages 261–264, New York, NY, USA, 2004. ACM Press.

[2] Mahoro Anabuki, Hiroyuki Kakuta, Hiroyuki Yamamoto, and Hideyuki Tamura. Welbo: an embodied conversational agent living in mixed reality space. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 10–11, New York, NY, USA, 2000. ACM Press.

[3] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. Digital tape drawing. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 161–169, New York, NY, USA, 1999. ACM Press.

[4] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. Digital tape drawing. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 161–169, New York, NY, USA, 1999. ACM Press.

[5] Ravin Balakrishnan and Ken Hinckley. The role of kinesthetic reference frames in two-handed input performance. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 171–178, New York, NY, USA, 1999. ACM Press.

[6] Robert M. Berne, Matthew N. Levy, Bruce M. Koeppen, and Bruce A. Stanton. *Physiology, Updated Edition: With STUDENT CONSULT Online Access (Physiology).* Mosby, 2004.

[7] W. M. Budney and S. K. Holewa. Midas: integrated cad for total system design. In *DAC '85: Proceedings of the 22nd ACM/IEEE conference on Design automation*, pages 529–535, New York, NY, USA, 1985. ACM Press.

[8] Andrew Carvey, Jim Gouldstone, Pallavi Vedurumudi, Adam Whiton, and Hiroshi Ishii. Rubber shark as user interface. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 634–639, New York, NY, USA, 2006. ACM Press.

[9] M. Fiorentino, G. Monno, and A. E. Uva. Smart tools for virtual reality based cad. In *Innovazione nella Progettazione Industriale*, pages 1–15. ADM-AIAS International Conference, 2004.

[10] George W. Fitzmaurice, Hiroshi Ishii, and William A. S. Buxton. Bricks: laying the foundations for graspable user interfaces. In *CHI '95: Proceedings of the SIGCHI conference*

on Human factors in computing systems, pages 442–449, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

[11] Neil Gershenfeld. *FAB: The Coming Revolution on Your Desktop–From Personal Computers to Personal Fabrication.* Basic Books, 2005.

[12] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. Creating principal 3d curves with digital tape drawing. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 121–128, New York, NY, USA, 2002. ACM Press.

[13] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. Creating principal 3d curves with digital tape drawing. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 121–128, New York, NY, USA, 2002. ACM Press.

[14] Ken Hinckley. Input technologies and techniques. pages 151–168, 2003.

[15] Ken Hinckley, Randy Pausch, Dennis Proffitt, and Neal F. Kassell. Two-handed virtual manipulation. *ACM Trans. Comput.-Hum. Interact.*, 5(3):260–302, 1998.

[16] Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassell. Cooperative bimanual action. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 27–34, New York, NY, USA, 1997. ACM Press.

[17] Neil Hopkinson, Richard Hague, and Philip Dickens. *Rapid Manufacturing: An Industrial Revolution for the Digital Age.* John Wiley and Sons, 2006.

[18] Hiroshi Ishii and Brygg Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 234–241, New York, NY, USA, 1997. ACM Press.

[19] H. Kato and M. Billinghurst. Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In *In Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*, 1999.

[20] Boriana Koleva, Holger Schnädelbach, Steve Benford, and Chris Greenhalgh. Traversable interfaces between real and virtual worlds. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 233–240, New York, NY, USA, 2000. ACM Press.

[21] Falko Kuester, Mark A. Duchaineau, Bernd Hamann, Kenneth I. Joy, and Antonio E. Uva. 3divs: 3-dimensional immersive virtual sculpting. In *NPIVM '99: Proceedings of the 1999 workshop on new paradigms in information visualization and manipulation in conjunction with the eighth ACM internation conference on Information and knowledge management*, pages 92–96, New York, NY, USA, 1999. ACM Press.

[22] Paul Lapides, Ehud Sharlin, Mario Costa Sousa, and Lisa Streit. The 3d tractus: A three-dimensional drawing board. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 169–176, Washington, DC, USA, 2006. IEEE Computer Society.

[23] Andrea Leganchuk, Shumin Zhai, and William Buxton. Manual and cognitive benefits of two-handed input: an experimental study. *ACM Trans. Comput.-Hum. Interact.*, 5(4):326–359, 1998.

[24] Andrea Leganchuk, Shumin Zhai, and William Buxton. Manual and cognitive benefits of two-handed input: an experimental study. *ACM Trans. Comput.-Hum. Interact.*, 5(4):326–359, 1998.

[25] Jeffrey K. Liker, Mitchell Fleischer, Mitsuo Nagamachi, and Michael S. Zonnevylle. Designers and their machines: Cad use and support in the us and japan. *Commun. ACM*, 35(2):77–95, 1992.

[26] Masatoshi Matsumiya, Haruo Takemura, and Naokazu Yokoya. An immersive modeling system for 3d free-form design using implicit surfaces. In *VRST '00: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 67–74, New York, NY, USA, 2000. ACM Press.

[27] Kevin T. McDonnell, Hong Qin, and Robert A. Wlodarczyk. Virtual clay: a real-time sculpting system with haptic toolkits. In *SI3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics*, pages 179–190, New York, NY, USA, 2001. ACM Press.

[28] JP Merlet. Direct kinematics and assembly modes of parallel manipulators. *International Journal of Robotics Research*, 11(2):150–162, 1992.

[29] P. Milgram and F. Kishino. A taxonomy of mixed reality visual displays. In *IEICE Trans. Inf. Syst.*, number E77-D, 1994.

[30] Mark R. Mine, Jr. Frederick P. Brooks, and Carlo H. Sequin. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 19–26, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[31] James Patten, Hiroshi Ishii, Jim Hines, and Gian Pangaro. Sensetable: a wireless object tracking platform for tangible user interfaces. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 253–260, New York, NY, USA, 2001. ACM Press.

[32] James McMichael Patten. *Mechanical constraints as common ground between people and computers.* PhD thesis, Massachusetts Institute of Technology, 2006.

[33] Ben Piper, Carlo Ratti, and Hiroshi Ishii. Illuminating clay: a 3-d tangible interface for landscape analysis. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 355–362, New York, NY, USA, 2002. ACM Press.

[34] Mengia S. Rioult-Pedotti, Daniel Friedman, Grzegorz Hess, and John P. Donoghue. Strengthening of horizontal cortical connections following skill learning. *Nature Neuroscience*, 1:230–234, 1998.

[35] M. Rohs and B. Gfeller. Using camera-equipped mobile phones for interacting with real-world objects, 2004.

[36] Emanuel Sachs, Andrew Roberts, and David Stoops. 3-draw: A tool for designing 3d shapes. *IEEE Comput. Graph. Appl.*, 11(6):18–26, 1991.

[37] Steven Schkolne, Michael Pruett, and Peter Schröder. Surface drawing: creating organic 3d shapes with the hand and tangible tools. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 261–268, New York, NY, USA, 2001. ACM Press.

[38] B Sener, PW Wormald, and RI Campbell. Towards"virtual clay" modelling - challenges and recommendations: A brief summary of the literature. In *Proceedings of the 7th International Design Conference*, pages 545–550. The Design Society, 2002.

[39] Ehud Sharlin, Benjamin Watson, Yoshifumi Kitamura, Fumio Kishino, and Yuichi Itoh. On tangible user interfaces, humans and spatiality. *Personal Ubiquitous Comput.*, 8(5):338–346, 2004.

[40] Laura Slack. *What is Product Design? (Essential Design Handbooks).* RotoVision, 2006.

[41] Michael Tsang, George W. Fitzmaurice, Gordon Kurtenbach, Azam Khan, and Bill Buxton. Boom chameleon: simultaneous capture of 3d viewpoint, voice and gesture annotations on a spatially-aware display. In *UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 111–120, New York, NY, USA, 2002. ACM Press.

[42] B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Syst. J.*, 39(3-4):915–931, 2000.

[43] Brygg Ullmer and Hiroshi Ishii. The metadesk: models and prototypes for tangible user interfaces. In *UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology*, pages 223–232, New York, NY, USA, 1997. ACM Press.

[44] Brygg Anders Ullmer. *Tangible Interfaces for Manipulating Aggregates of Digital Information.* PhD thesis, Massachusetts Institute of Technology, 2002.

[45] John Underkoffler and Hiroshi Ishii. Urp: a luminous-tangible workbench for urban planning and design. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 386–393, New York, NY, USA, 1999. ACM Press.

[46] Sidney W. Wang and Arie E. Kaufman. Volume sculpting. In *SI3D '95: Proceedings of the 1995 symposium on Interactive 3D graphics*, pages 151–ff., New York, NY, USA, 1995. ACM Press.

[47] Stephen H. Westin. Computer-aided industrial design. *SIGGRAPH Comput. Graph.*, 32(1):49–52, 1998.

[48] Christopher D. Wickens, John D. Lee, Yili Liu, and Sallie Gordon-Becker. *Introduction to Human Factors Engineering (2nd Edition).* Prentice Hall, 2003.

7387 - 37