An Emergent Framework for Designers working in Physical/Digital Spaces

David Gullick

MSc Digital Innovations (Lancaster University)

Digital Innovation

School of Computing and Communications

Lancaster University

March 22, 2018

This thesis is submitted in partial fulfilment of the requirements

for the degree of Doctor of Philosophy

Abstract

Technology is becoming more deeply entwined with the spaces in which we live every day. As it does so, the line that divides that which is considered digital, and that which is physical is becoming blurred. As these two spaces merge, the elements that contribute to the way in which we understand to interact within them become harder to define. The work described within this thesis focuses on exploring this space using a formalised methodology that mirrors the design process over a number of iterative and exploratory "Research through Design" projects. This work highlights and discusses a number of key themes that reoccur throughout these projects, and then augments an established interaction design framework to incorporate these themes. Finally, reflections on this formalised design process, and the future of this hybrid space are discussed.

Acknowledgements

Firstly, I would like to express my sincere gratitude to my adviser Prof. Paul Coulton for the continuous support of my Ph.D study and related research. This journey into the world of design would not have been the same without your guidance. I would also extend my thanks to Dr Manfred Lau for his guidance throughout this project.

I would like to thank those at the Highwire Centre and the Creative eXchange: it was there that I had my first exposure to 'designerly way of thinking' and the freedom to explore concepts which I would have otherwise been unable.

Last but not the least, I would like to thank my family, and particularly my partner, Lucy, for supporting me throughout writing this thesis. I would not have finished it without you.

This work has been funded by the Digital Economy programme (RCUK Grant EP/G037582/1), which supports the Highwire Centre for Doctoral Training at Lancaster University.

Declaration

This thesis is my own work and no portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification at this or any other institute of learning.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Preface - About the author and setting of the scene

This thesis is the result of four years of research work carried out in the Highwire Doctoral Training Centre at Lancaster University, and was funded by a grant from the Engineering and Physical Sciences Research Council (EPSRC) as part of a doctoral training initiative. The criteria for this grant includes a heavy emphasis on cross disciplinary work - the academic equivalent of pushing potential candidates to work in realms beyond their “comfort zone”. My background is in Computer Science, and as such, I had a strong record of building things and then assessing them (in what I would now call) a '“Computer Science” way. Since being coerced out of my “comfort zone” by the PhD programme, I have found a keen interest in the interaction that takes place when a person uses a piece of technology, rather than focusing on the technology itself. This interaction, the “hows” and “whys”, and causes and effects seemed less immediately obvious to me than the intrinsic operation of the technology involved. The way I had measured success in technology up to this points was now inadequate. For example qualities such as 'fun' were not apparent when comparing the time it took people to complete two tasks. As a result, I turned to alternative approaches, such as Design, for help - to find that a more qualitative approach is not just acceptable, but in fact the norm. As a result, this work is being presented in the format of a design thesis, but with a heavy influence from Computer Science - Computer Science can be considered the “how” behind this work, whilst interaction design research can be considered the “why”. I will go further into the reasoning behind this in Chapter 6, but suffice to say that my journey throughout the PhD process has allowed me the lateral academic movement that I needed to define myself as a researcher. Without this personal academic development,I would not have had the freedom to approach problems, such as those described in this thesis, from a design perspective.

Whilst my PhD programme is described above as having the goal of being cross discipline, the majority of my time throughout this period of research was spent in a nearby design office, which was part of what is called Creative eXchange (CX). CX is a design research group which acts as a collective effort between Lancaster University, The Royal College of Art and Newcastle University, to work on improving links between the creative industry and academics. Whilst this arguably has little bearing on this work in an official capacity, working within this space allowed me extended insight into the design process, and the way that research is carried out in CX. I feel that the project-focused and iterative research cycle adopted by those in CX has greatly influenced the style and direction of this work.

1.2 Introduction

Throughout the last sixty years, continuous advances in computing technology have meant that computers are now an integral part of life for many of us. Computing technology has been improving in an almost exponential fashion (Moore et al., 1998), in almost every respect including: speed, storage capacity, physical size, connectivity, and network throughput. In fact the way in which people exist alongside, and interact with, computers has changed. In most cases, people no longer have a many-to-one relationship with computers (e.g. old punch card computers like the Electronic Integrator and Computer (ENIAC) of the 1940s, where people rented time on one computer), but rather have a one-to-many relationship with computing technology. Whether or not we use a computer at work or at home, we probably use multiple computers in some form or another: our mobile phone, car, microwave, or one of many other devices with which we interact in our daily lives. Due to this ubiquity, the line that separates the digital from the non digital is becoming more and more ambiguous. With the rise of the "hacker" mindset, the Internet of Things (IoT), and the increasing availability of cheap, yet powerful, technology, the concept that we can clearly label and define what is physical and what is digital is becoming obsolete. However, whilst the technology that is allowing this change is interesting and worthy of much research in its own right, it is not the focus of the discussion contained within this thesis. Instead, this work focuses on gaining some insight in to the way we interact with computer technology, particularly in spaces that contain both physical and digital elements. Furthermore, how do we design for these interactions given that the spaces that these interactions are taking place are themselves changing and evolving as new technology is incorporated? This thesis hopes to highlight some concepts that should be considered by designers when working within this hybrid space. The end goal here is not to produce “rules” for designers - the concept of laws seems non-applicable to the subjective nature of the field of design - instead, the aim is to highlight areas for consideration that may have not previously been raised.

As a reader, you might question the need to treat these spaces that consist of both physical and digital elements differently than a purely physical space. Do we need to consider both the physical and the digital? Why are these guidelines necessary?

To begin to frame an answer to these questions, I have chosen the following anecdote shared by one of the participants a study describes later on in this work. The participant discusses the way that airports have changed in atmosphere since digital screens have been used to display boarding times. “In the ‘good old days’”, airports used mechanical “split-flap” display boards. When these boards change, there is a loud ticking noise as the split-flaps physically rotate due to their inner mechanical workings. The result of this is that people only looked at the boards when they heard the mechanical noise once every few minutes, which signalled that there had been an update to the screen. Consequently people would spend the rest of their time engaged with something else, such as reading a book, or browsing in a shop. In most airports these boards have now been replaced by large Liquid Crystal Display (LCD) displays, which do not produce any audible sign that the information is changing. The participant described the way that modern airway travellers now spend most of their time in airports looking up at these boards waiting for it to change.

In this scenario, the departure times’ display board has been replaced with one that is far superior in terms of visually displaying information, yet it seems that the audible cue that new information is available was either purposefully ignored, or overlooked by the designer as a key part of the interaction between travellers and the technology. As a result, the interaction with the replacement board is very different to the interaction with the old piece of technology. In reference to our goal, it is not the fact that designer has not included an audible cue that we want to change, rather we want to highlight how important it is that the audible cue was a consideration in the design of the interaction, alongside the other physical attributes of the original display. Acknowledging these physical properties of technology is important for two clear reasons: the physical embodiment of the technology results in it being perceivable in ways that may not have been key to its main function; and interactions that take place between human and computer can often include, or rely on, these more physical elements to function.

Although the airport display board scenario is, arguably, on the periphery of the design space that we are discussing, it contains both physical and digital elements. Even the small changed discussed above had a profound effect on those people using the airport. It seems that if we are going to continue to design for these increasingly common "hybrid" spaces, we will need to consider both the physical context of the space, as well as the virtual, and also have a more formal way to classify components of each of them.

These kind of trade-off between the physical and the virtual elements within a space is not only apparent in the above air port scenario but also in many other recent technologies. For example, early touch screen phones replaced the physical keys with on screen buttons. This process means the haptic and tactile feedback given to us by the hardware of a button is lost - that is the "feel" of the button as we press it is no longer evident. It has been suggested in relation to touch screen devices (Hoggan et al., 2008) that, where possible, haptic feedback should be given (possibly via vibration) where no physical buttons exist to maximise their usability. This mirrors my personal experience: it was only once mobile phones began to offer a vibrate function to indicate that the button had been pressed that on screen keyboard became acceptable to the majority of users. Similarly, projection based keyboard technologies (that project a keyboard onto an otherwise flat, non interactive surface), whilst space saving, offer no haptic feedback and did not have the uptake that some expected. There may been other factors at play, but for many, myself included, the lack of keys and lack of tactile feedback was responsible for its unpopularity.

Given these scenarios as examples, it seems clear that much of Human Computer Interaction (HCI) is about the physical as well as the digital though it is not always considered as such. I am not arguing that the move of systems from entirely physical to entirely digital, which is typical of today, is a bad one at all. In most cases, it is done well, and results in a system that is more reliable and more efficient. In fact, this move is almost inevitable given the current trend of incorporating technology into our everyday lives. As the spaces in which we exist are becoming hybrid physical/digital spaces, I am arguing for the more thoughtful inclusion and/or consideration of this hybrid nature throughout the design process - I believe that the interaction should be the focus of the design, rather than the technology.

Rather than trying to distill this problem into a single and well defined question, I have chosen to begin this research with an investigation into a problem area. The reasons for this are discussed in detail in Chapter 2. In brief, the reduction of this problem to a single question removes much of the context for its exploration. At this early stage, this seems counter-productive. Interactions and experience can be greatly changed by even a small detail. If we look back to the the airport display anecdote, we can see that even the absence of one sound in the move between a mechanical screen to a digital one had a great effect on the experience of travellers. This is because interaction and experience is dependant not only on the object but also on the surrounding context within which this interaction occurs. A reductionist approach would likely remove much of the context which would affect the observed interaction behaviour. A reduction of this problem space also leaves the question of exactly *how* it is reduced. Given that this area is still largely unexplored, exact questions regarding this space are yet to be formulated by the research community. Concentrating on a particular question at this early stage means any subsequent research will be viewed through a singular corresponding lens focused on finding data related to that question alone. If, during the research, other patterns emerge, these are likely to be missed as they may be outside of the scope of our lens. More so, this lens will inherently affect my actions as a researcher, such as my decisions on data collection methods. Whilst certain methods may be suited to answer a distilled question, they may not collect sufficient data to allow further themes to emerge. In short, with a reductionist question rather than a wider exploration of a problem area we are likely to ignore the effects of context and allow our own biases and understanding to limit and narrow our research scope.

As an alternative to the reductionist approach, I observed and inductively analysed a number of research projects that exist within this problem space for emergent themes. Using triangulation of these themes, I produce some concepts and guidelines that are applicable to the problem space in general. This research takes the form of multiple projects that exist within an iterative research cycle which is described in detail in Chapter 2, with relevant supporting literature. The remainder of this chapter acts as an introduction to core concepts related to working in this area - each concept can be considered key for understanding, and mentally framing, the remainder of the work discussed in this thesis. Whilst these concepts are being introduced at the beginning of this thesis, they were discovered somewhat organically throughout the iterative research process discussed in Chapter 2. As such, they are being introduced here as a set of cohesive ideas that will supply a foundation for the remaining work. Where appropriate, each discussed project will introduce more specific theories.

1.3 Affordances

In order to further understand the importance of the physical in our interactions, we must first discuss how people understand how to interact with their surroundings.Gibson (1977) introduced the concept of "affordances" as properties of an object that lend themselves to a particular action or "action possibilities" that exist in the environment, regardless of whether that action possibility is acted upon or not.

"what we perceive when we look at objects are their affordances, not their qualities . . . what the object affords us is what we normally pay attention to" (Gibson, 1977, p. 134)

In saying this he voices the view that objects are not inherently specific about what they are and how to interact with them, these things are instead ascribed by the viewer. Whilst Gibson's definition of affordances is grounded in the field of cognitive psychology, Norman (1999, 2002) transposes this into the context of HCI by coining the term "perceived affordance". Perceived affordance is the affordance of an object given the context of the system: including the influence of societal norms upon the users understanding and use of a object; instinctual associations the user may make; and the user's current understanding and expectations of interaction with the object. Affordances are different for each person, dependant on their personal context. If we use computerised displays as an example, we can compare usage of displays from the past and today. In the past, most screens were not touch screen enabled. As a result, most people would not have considered touching the screen for information. However, more recently, touch screens have become more prevalent, shifting the general understanding of a screen from non-touchable to touchable. From now on, most people will assume that a screen could have touch screen properties as a direct result of the cultural and technological change in society.

Similarly Gaver (1991) divided affordances into three categories - "perceptible", "hidden", and "false". Perceptible affordances exist when detectable information for an existing affordance is apparent to the user, leading the user to perceive the affordance and act upon it (e.g if we see a button, we might think we can press it). In comparison, hidden affordances still exist, but do not offer any perceptible information to the user regarding their existence (e.g responsive to touch, but there is no button). A false affordance exists when an apparent affordance has no real function (e.g a button that cannot be pressed). To apply these terms to the previous example of a user touching a non touch-screen display, this action would be based upon a false affordance if the user believed it was touch screen responsive. Gaver proceeded to divide perceptible affordances into two sub-categories. The first of these, "sequential" affordances, are affordances that reveal themselves over time:

"Affordances are not passively perceived, but explored. Learning is a matter of attention rather than inference." Gaver (1991, p. 82)

Prolonged or successful interaction may provide new information that reveals a previously hidden affordance to the user. The second of Gaver's subcategories are "nested" affordances. Nested affordances are similarly located, and knowledge of one improves the understanding of another (i.e objects locate near each other may indicate how each other can be used). Gaver also introduces the term "Correct Rejection", which is appropriate when neither an affordance nor any misperceptions of the aforementioned affordance exist (e.g the user believes the screen is not touch screen, and it isn't: they have correctly rejected touch). These categories have been included in Figure 1.1.

Whilst both Norman and Gaver are recognisable contributors, McGrenere and Ho (2000) advance previous works by suggesting that to regard affordances as binary is to oversimplify them. They suggest that affordances should be considered as existing within a two-dimensional space, between clarity of information (degree of perceptual Information) and ease of affordance (degree of affordance). Interaction designers should aim to increase the position of the interaction on both of these axes.

A good way to highlight the importance of the physical elements and affordances of an object in an interaction is to consider a scenario where the speed and efficiency of interactions are critical. Sticking with the aircraft theme, MacKay (1999) in the paper "Is paper safer? The role of paper flight strips in air traffic control", describes a project in 1999 that attempted to revolutionise air traffic control systems by moving from a paper format to a digital one (also

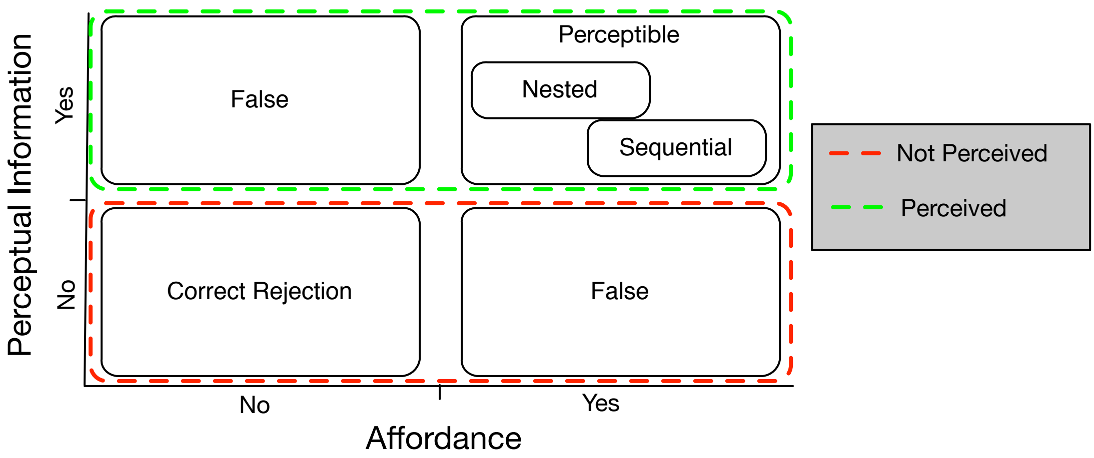

Figure 1.1: A rework of Gaver's original affordances diagram, augmented with Nested and Sequential affordances.

discussed by Harper et al. (1990), Hopkin (1993), Preux (1994)). At the time, the system for air traffic controllers was for each controller to use a number of paper-strips, one per air craft, manoeuvred manually on a surface. The act of controlling these aircraft sufficiently was inherently "multi-threaded"; Mackay described how each controller spent approximately two seconds looking at each strip of paper, considering the referenced aircraft's altitude and heading alongside other parameters before moving on to the next piece of paper. Instead of just being a reference, these pieces of paper also performed a number of other functions: if the controller needed to mark something about an aircraft, they could easily move the piece of paper to the side, if they needed to differentiate further, they could fold a corner; to change the order of the aircraft, they could slide the paper strips around the desk; they could record extra information on them using a pen; if the air traffic controller agreed something with the aircraft (such as a new altitude), the controller could sign the paper, and it would act as a contract. The mere fact that this strip was a physical object meant that it needed to be physically removed from a printer, mentally registered by the person printing it, and placed. It gave the user "the sense of 'owning' the aircraft and reinforces their memory of the current situation" (MacKay, 1999, p. 322).

Interestingly, when Mackay and the team of researchers describe how they attempted to replace this system with a digital alternative, they endeavoured to take much of their observations about the actions of air traffic controllers into account. However, they found that this alternative was ultimately unsuccessful, and that the new system would require a radical change in the existing work practises of the controllers. The new system lacked many of the affordances offered by paper which ultimately cost controllers time and accuracy which was unacceptable to them. In fact, elements such as the ambient noise level in the control room which the controllers used to gauge a sense of the workload were also affected as a result of the change. This is not to say that the digital system did not afford itself to the correct actions, but shows that the controllers were more experienced with physical interactions and so were quicker to perceive physical affordances than digital affordances and act upon them. Eventually, Mackay suggested keeping the physical strips of paper for the controllers to work with, but having a computer system that could detect and read these strips of paper as a viable alternative.

1.4 Tangible User Interfaces

The air traffic control room scenario is a perfect example of why it is important that our experience of interacting with the physical word is acknowledged when designing. Many other researchers are in agreement upon this point, and a whole field of research exists where understanding, and working with, our expertise from the physical world with digital systems is the goal: this is the field of Tangible User Interface (TUI).

Ishii and Ullmer (1997) presented TUI (originally called "Graspable User Interfaces" (Fitzmaurice et al., 1995)) as an alternative interaction method to Graphical User Interface (GUI) that was typical of HCI at the time. Building upon traditional ideas from Ubiquitous Computing (UBICOMP) and Augmented Reality (AR) research, Ishii and Ullmer suggest that "although we have developed significant skills and work practises for processing information through haptic interactions with physical objects ... most of these practises are neglected in current HCI design" (Ishii and Ullmer, 1997, p. 1). They suggested a tangible approach in which digital information is situated in physical space, and can be manipulated and perceived as a physical object, therefore benefiting from our experiences of, and inheriting the affordances of, an object in the physical world. A good example of this concept is musicBottles (Ishii et al., 2001) in which a system is built that allows the user to "store" music inside glass bottles. When the lid is removed from a glass bottle the music stored inside is played. The argument for this method of interaction is made when we consider that most people have already interacted with a glass bottle before and are therefore already experts in this kind of interaction; there is no need to learn a new interface as the user is already comfortable interacting with the physical world interface. Obviously no music is actually stored within the bottle, but a system is in place to act as if it is. In this case, augmentations to the interactions should be easy for the user to understand as they follow the same rules as the physical world. Some examples are shown in Figure 1.2.

Aim	Solution
Start playing music.	Open lid.
Stop playing music.	Close lid.
Increase volume.	Open lid more.
Decrease volume.	Close lid more.
Play multiple sounds. at once.	Open multiple lids.
Mix multiple sounds.	Pour one bottle into another.
Delete sound.	Pour music away.

Figure 1.2: Example aims and solutions of a musicBottles like system.

The field of TUI has produced a number of relevant and interesting ideas. For example, Ishii uses the term “Ambient media” (Ishii and Ullmer, 1997, Ishii et al., 1998) to discuss technology that communicates to the user in a non-intrusive “ambient” form. Consider the scenario of an office worker at a desk. They might know that someone is waiting to meet with them as they can perceive them waiting outside the office door though their peripheral vision. They can tell that the weather outside has become worse as the level of light coming through the window has dropped, and they can hear raindrops. Similarly, in the aforementioned air traffic control scenario, the ambient noise level in the room passively informs the controllers about the level of work in the room. Ambient media communicates with the user in a similar fashion. It does not grab your attention, but passively provides information.

1.5 Natural User Interfaces

Natural User Interface (NUI) as a research area is similar in ideals to TUI: its key aims are to “allow the users to act and feel natural” (Wigdor and Wixon, 2011) in their interactions with technologies. However, it does not aim to make the interactions possible without some kind of learning (e.g by utilising the skills and practises that people have developed for the real world as with TUI), but rather it aims to make the the interaction feel as natural as possible to users. Whilst it may require some learning to advance from a novice user to an expert user, the interaction should feel equally natural to all. For example the input languages for the Palm

Pilot Graffiti and Graffiti 2 (Költringer and Grechenig, 2004) have similarities to the English language and use an interaction with which most people will feel comfortable : the same basic interaction as pen on paper. Whilst experts will learn the nuances of the software over time, the interaction will be natural to everyone comfortable with a pen.

1.6 Tabletop Interfaces

Tabletop Interfaces in the context of this work can be considered a subset of TUI that specifically focus on interactions taking place in a tabletop scenario. A prominent example of such projects is "ReacTable" (Jordà et al., 2007) which explores the interaction between tabletop surface and tangible objects as controllers for the purpose of audio technology applications.

In such projects the tabletop usually takes an active roll in the interaction itself. In most cases the table is used as a receiver for projected imagery (e.g Patten and Ishii (2007), Underkoffler and Ishii (1999), Waldner et al. (2006), some also allow touch or gestures (Schöning et al., 2010), some focus on object tracking (Jordà et al., 2007) and others even modify their shape as part of interactions(Follmer et al., 2013).

As interactions explored later within this thesis take place in a similar tabletop context (and if not a tabletop then a flat surface). These interactions can therefore be considered as related. Although relevent, the aim of this work is directed more at TUI in general and should not be considered as only applicable to the tabletop context.

1.7 How do we Do? How do we Know? How do we Feel?

The concept of affordance and the goals held by the TUI and NUI research community resonate with me as a researcher, and support my argument that the physicality of an interaction should be considered during the design process. To further investigate this realm of physical/digital interaction, we also need to consider how a person understands how to interact with a whole system, rather than an individual object. Whilst there are a number of frameworks in the relevant literature that help conceptualise this, for this work I am using the framing that Verplank describes in his Interaction Design Sketchbook (Verplank, 2009). Verplanks framework for modelling interactions has been chosen for three reasons: firstly, it offers abstraction over the details of particular interaction techniques. This means it does not focus on the particular interaction that is being carried out, but offers a higher level framing for the process. For example, the field of Tangible User Interfaces has multiple frameworks within relevant literature (such as Ullmer and Ishii (2000)), but each very specific to one area of research. Verplanks framework manages to encompass wider variety of interactions by looking at the underlying reasons behind an interactions success. Secondly, due to its simplicity, it is particularly easy to place interactions within Verplank's framework. Lastly, Verplanks framework manages to capture an aspect of interaction often overlooked by other frameworks; the type of media with which the user is interacting (this is explained in more detail in Section 1.7.2).

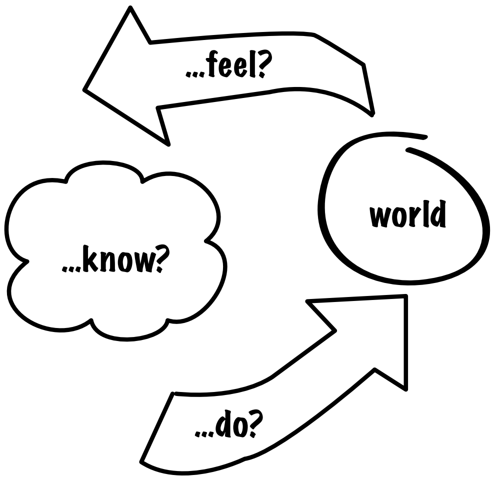

Figure 1.3: An illustration showing Verplank's do, know, feel cycle.

Verplank approaches the problem of interaction design as three questions: How do you do? How do you feel? and How do you know?: " I can flip [do] a light switch and see (feel?) the light come on; what I need to know is the mapping from switch to light"(Verplank, 2009, p.6)

The greater the distance, or abstraction, between the input and the output, the more varied the user's conceptual model may be. These questions describe the thought process that occurs when someone interacts with an object. A person should be able to understand the current state of the system, and affect it by introducing change. This change will then be learned as a cause and effect rule. Figure 1.3 shows this as a cycle that may repeat as the user interacts with the system.

1.7.1 Discrete vs Continuous

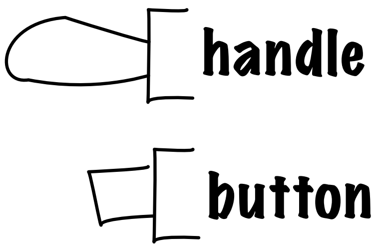

Figure 1.4: An illustration showing Verplank's Differentiation of Discrete and Continuous

The example above consists of a light and a switch (button), an input that consists of only two states; ON and OFF. This type of input can therefore be considered discrete as there is no in-between state. Other inputs such as handles may have not have a discrete number of states and are therefore more continuous in nature. Verplanks framework labels interactions (or DOing actions) as either discrete or continuous(see Figure 1.4).

1.7.2 Hot vs Cool Media

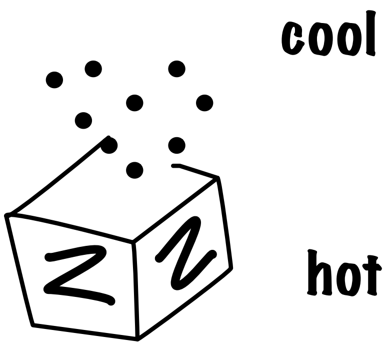

Figure 1.5: An illustration showing Verplank's concepts of Hot vs Cool media.

According to McLuhan (1994), all media can be categorised as either Hot or Cool, and Verplank uses these labels in his description of media types. Hot media demands our attention or dominates one particular sense, leaving little room for participation. Cool media in comparison does the opposite, it is usually spread across multiple senses and actively invites participation. If we apply this to a GUI based desktop computer interaction scenario, we can imagine that something such as a Portable Document Format (PDF) document can be considered a Hot media (it is rigid and demands your attention), whereas a video game can be considered Cool media (spread across multiple senses, invites interaction, and can be less demanding of attention depending on the game).

1.7.3 Map & Path

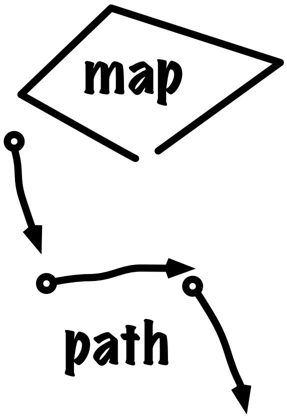

Figure 1.6: An illustration showing Verplank's concepts of Map and Path to knowledge.

Verplank's framework describes two different types of knowledge based upon the work by Lynch (1960). The first of these, "path knowledge", only requires knowledge of instructions to achieve a goal, and is most applicable to first-time users who will expect step-by-step instructions on the correct course of action to follow. The second type of knowledge, "map knowledge", is based upon Lynchs work on "imageability", referring to the existence of a mental map in the users mind that allows them to independently generate a path of instructions to solve a new problem.

1.8 Interaction Space

Given this overview of the way in which people can be seen to understand how to interact with objects and the world around them, we can begin to more clearly define the problem space that we are investigating. Until relatively recently, there has always been an obvious divide between what is considered digital and what is considered physical. This divide is typically a screen of some sort. A person exists in the physical world, and looks though some window interface (a screen) to observe and interact with a digital one (e.g playing a video game or browsing the Internet). As previously discussed, due to advances in technology, and its ubiquity, this is no longer necessarily the case. In reference to video games, Juul (2010) describes this space as being divided into three distinct areas “Player Space”, “Screen Space”, and “3D Space”. One possible configuration of these spaces is shown in Figure 1.7.

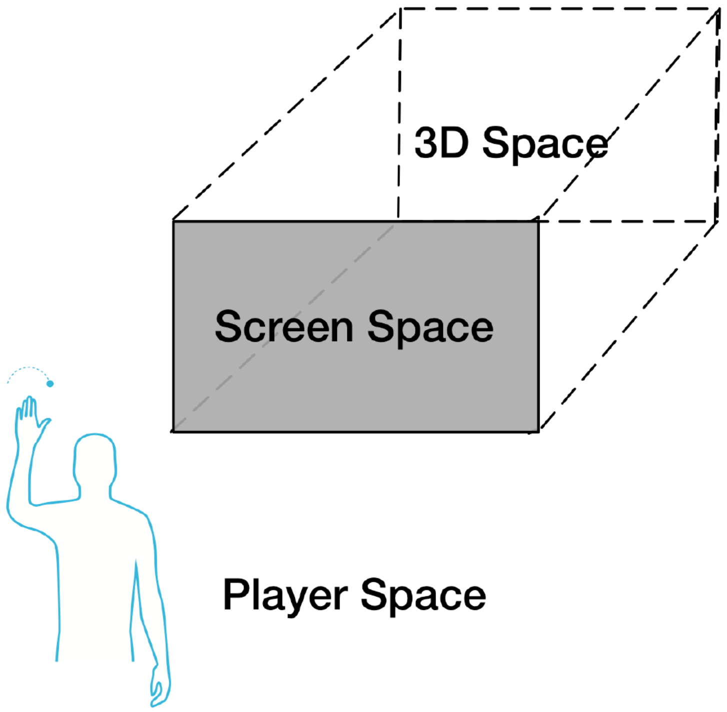

Figure 1.7: An example of Juul’s division of game interaction space.

This division of game space provides a useful way for designers to consider where the focus of attention for the player might be when interacting game objects in screen-based scenarios. This consideration will allow designers to clarify in which space, and how, the interaction takes place and in which space, and how, feedback on that interaction is presented to the player. Whilst we are not explicitly focused on dealing with games, Juul’s game space terminology offers a formalisation of spaces and a clearer description of the space we are investigating. Given the typical arrangement of game spaces as shown in Figure 1.7, we can represent the “hybrid” physical/digital spaces as shown in Figure 1.8.

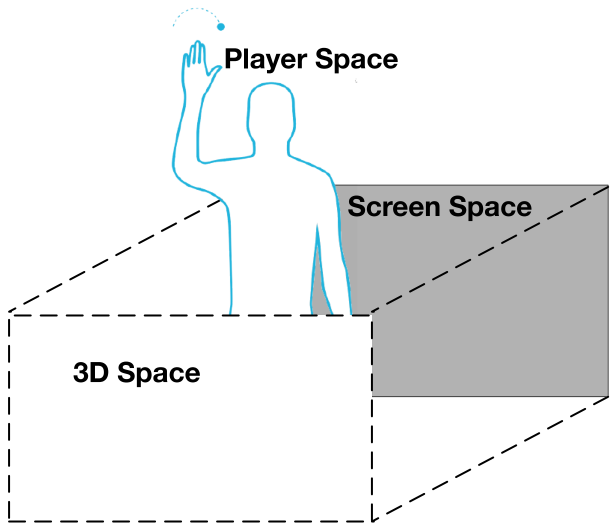

Figure 1.8: The space I am investigating, described using Juul's terminology.

1.9 Liminal and Liminoid spaces

Figure 1.8 shows the player now positioned within the 3D space within which they are interacting. This space can be considered to now exist as a hybrid space with both physical, and virtual elements. This space can be seen as existing as a reality as a point along the "virtuality continuum": a continuum of reality spaces that exists between two endpoints - the entirely real, and the completely virtual (Milgram et al., 1995). Within this continuum exist all other hybrid realities, with their ratio of real/virtual dictating at which point they stand (see Figure 1.9). Historically, users have existed at one of of this continuum (real), and have transitioned to another place on this continuum by using a Virtual Reality (VR) headset, or interacting with some kind of technology. It could be argued that this puts users in a permanently transitional state, forever at the threshold between the real and the virtual (Coulton, 2017). However, more recently technology has evolved to the point that we no longer exist at the threshold between the real and the virtual; instead the virtual space can be seen to have grown so that it overlaps portions of the real - placing users in a liminoid state, a state which is simultaneously real and virtual (see Figure 1.10).

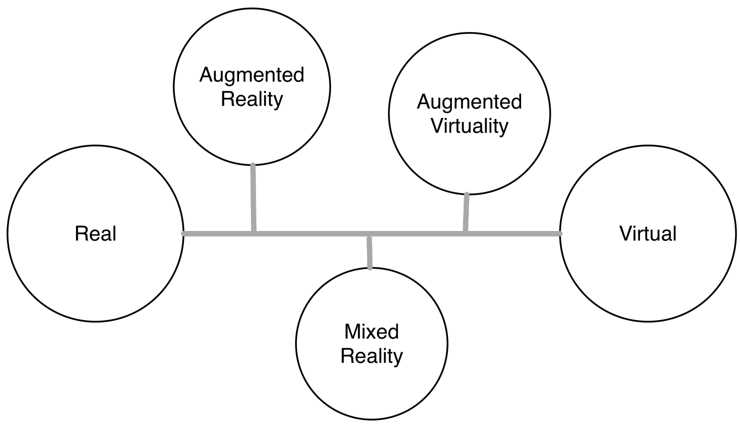

Figure 1.9: An illustration of the Virtuality Continuum

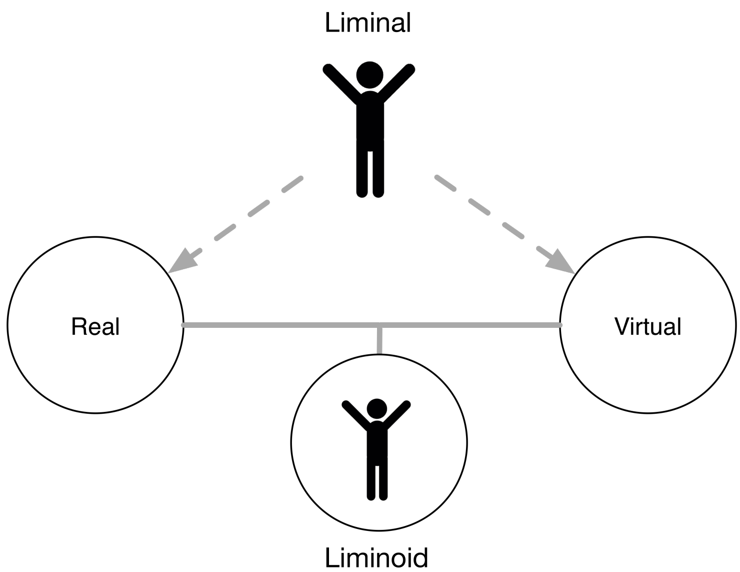

Figure 1.10: Liminal vs Liminoid space

In fact, it has become impractical to differentiate these space with the terms like "real" and "virtual", the "physical" or the "digital" as these are now becoming more or less synonymous in certain spaces. A more apt label could be "Atoms" and "Bits" - Atoms labels the real, or the physical, and Bits labels the virtual or the digital aspect. Whilst it can be argued that Bits also exists as atoms, I consider Bits in this context to refer to computer code.

1.10 Internet Of Things (IoT)

The phrase "IoT" is used to describe the inter-networking of physical devices that allow such devices to exchange data. Typically, but not always, this takes the form of embedded and network capable sensors within physical objects. For example, the Nest thermostat[1] can sense the temperature, and humidity, of the space in which it is located, as well as detecting movement. With this information it can predict at what times it is most efficient to turn on the heating, and then communicates this to the boiler over the Internet. Similarly, Do It Yourself (DIY) IoT objects are becoming much more common given the cheap availability of modular, beginner friendly hardware such as the Arduino[2] (a small microprocessor) or Raspberry Pi[3]. Any object can now have digital properties or functions - even plant pots(ParrotPot) and kettles[4]. The relevance of the IoT to this work is the way in which these networked objects are further blurring the boundary between what is physical and what is digital.

1.11 Phygital Objects

"Phygital" Objects are closely related to objects that are part of the IoT; tshe difference lies within the interaction. IoT objects and surrounding research is typically focused around the objects themselves, i.e the physical. In contrast the term "Phygital" refers to similar objects but with the equal focus on the physical and virtual elements of their interactions. An example of such an object would be Activisions Skylanders[5] which exists as a video game with accompanying physical objects. The Skylanders objects can be used in conjunction with the video game and act as an embodiment of virtual attributes which can be adjusted and transferred with the game object. These objects are of interest to this work as they, in a similar fashion to IoT objects, exist within spaces both physically and virtually.

[1] https://nest.com/uk/thermostat/meet-nest-thermostat/
[2] https://www.arduino.cc/
[3] https://www.raspberrypi.org/
[4] https://www.amazon.co.uk/Smarter-iKettle-Kettle-Stainless-Steel/dp/B0179NGBT4
[5] https://www.skylanders.com/uk/en

1.12 Frameworks and Conclusions

As previously discussed rather than posing a specific question or hypothesis at the start of the research I present an aim to address a particular problem space through which answers to specific research questions can be found. It is through the studies and investigations described in this thesis, that this problem space becomes more focused and specific, evolving and changing as new knowledge is gained. As this new knowledge as accrued, and a clearer conceptual framing emerges it is easier and more appropriate to begin asking specific questions. As such, the overarching aim for this work can be considered the creation and/or description of a framework that supports work in this space.

Frameworks within research can exist in many forms, as the term in general described a basic supporting structure. Depending on the specific scenario this structure can be physical, such as is used in engineering; Digital, such as software frameworks the are often used by developers; or even conceptual such as those that encapsulate a problem or idea in such a way that it allows discussion and promotes conceptualisation. In this context however, the word framework should be considered as the latter of these three examples; a conceptual framework. Such frameworks act as a (or set of) conceptual guidelines that facilitate understanding by contextualising knowledge, and as such allow more specific questions to be posed.

Conceptual frameworks of this kind are prevalent in fast evolving or immature areas of research, such as computer science and TUI (e.g. Hornecker and Buur (2006), Koleva et al. (2003), Shaer and Hornecker (2010), Shaer et al. (2004), Ullmer and Ishii (2000)). For some relevant examples, consider Ullmer and Ishii (2000) discussing a conceptual framework built upon the 'MCRpd' interaction model for tangible interfaces - a model that considers the physical and digital representations of objects in a system and their uses in such systems. They argue that this conceptual tool acts as foundation for identifying and discussing several key characteristics of tangible user interfaces. Similarly, Hornecker and Buur (2006) discusses a conceptual framework for describing and denoting interactions in this space by building upon synthesising prior definitions into one conceptual framework. As such, the produced conceptual framework allows the discussion and exploration of interaction case studies in a way that was not easily performed before. As is typical in research, frameworks such as these are often produced in an iterative manner - as more knowledge (and technological advancement) is gained within a space, a conceptual framework is adjusted to accommodate.

As the space in which we are exploring is both immature (the field of TUI was only conceived in the later 1990's), and fast paced (much of the interactions described in the literature are driven by advancements in technology) a conceptual framework is suitable. Therefore the aims and objectives can be considered as follows:

Aim:

To produce a framework that contributes to an improved ability to conceptualise Mixed Physical/Digital spaces.

Objectives:

- Investigate and explain the epistemological standing for creating a framework.
- Explore the problem space considering the aforementioned epistemological standing through Research *through* Design.
- Interpret the findings from the exploration into a useable framework.
- Provide reflections upon this process, and the produced framework.

Chapter 2

Methodology

2.1 Introduction

In this chapter I discuss the positioning of this research with respect to what is considered valid knowledge, and what research methods best suit the acquisition of this knowledge. I have decided to approach this research with a methodology that better reflects the design process than some of the more traditional positivist methodologies used in HCI research. This is due partly to the open-ended nature of the research question, and partly to my epistemological positioning as a researcher. I believe that this decision can be best portrayed by considering this question posed by John Law:

"If this [something] is an awful mess . . . then would something less messy make a mess of describing it? . . . Simplicity . . . won't help us to understand mess"(Law, 2007, p. 3)

His discussion is one centred on a comparison of contemporary scientific techniques that favour clarity, specificity and repeatability at the cost of repressing "Mess" with an alternative view that attempts to embrace it. "Mess", according to Law. is almost the opposite of intellectual hygiene as is sought by many areas of research (e.g Weber (2015)) - According to Law, "Mess" can be considered everything that is typically removed in order to perform "lab-based" tests in more positivist approaches. He argues that this "Mess" makes up a very large portion of the world in which we exist, and as a result it is completely relevant to the research in terms of understanding both the data, and in fact defining (or defining) the question researchers are attempting to answer. When discussing the reasoning behind his view, Law uses anecdotal evidence to describe the difficulty of measuring, or indeed defining what to measure, when the research is taken outside of the "cleansed" lab environment. Real world examples used by Law include: complex social networks that are not always clear cut, or, the difficulty in healthcare of plotting the trajectory of the "typical patient". The difficulties with with research are now somewhat similar to the "Wicked Problems" described by Rittel and Webber (1973). Wicked problems are problems which are difficult to solve because of incomplete, contradictory and changing requirements which are not immediately obvious to the problem solver, and, in many cases, the researcher is not able to explicitly pinpoint the exact problem in the first instance.

I agree with Law when he talks about "Mess", and I believe the concept of HCI is moving in a direction that reinforce this viewpoint. When it was first conceived in the early 1970s as a field, the goal of HCI was to increase the usability of computers through a combination of engineering and cognitive science. The goal was not only to build systems that were more usable to people, but also to begin to model and understand why people interact with these systems as they do (Ghaoui, 2005, chap. 2). Since the 1970s, technology has improved and evolved considerably: for example Douglas Engelbart's invention of the computer mouse (English et al., 1967), the invention of the GUI, the Windows Icons Mouse Pointer (WIMP) , and the "Messy Desktop" metaphor have allowed even the least technically proficient users to increase productivity using computers (Reimer, 2005, Thacker et al., 1979). It is not just the processing power of computing devices that are increasing, but also their ubiquity is also increasing (UBICOMP(Weiser, 1994), the IoT (Ashton, 2009)). As a result, the original and rather naïve HCI definition of usability has continuously been re-articulated and re-defined: it now subsumes qualities such as: fun; well-being; collective efficacy; aesthetic tension; enhanced creativity; flow; and support for human development among others. What was once a simple question of increasing the usability of a system has become substantially more difficult with the addition of these considerations, and in many cases, it is now trickier to define an initial question for the researcher to solve. It seems obvious that in order to further understand how a person interacts with a system, we would need to consider things such as fun even though this attribute might be difficult to quantify; an interaction that is fast but people find non-enjoyable will probably be less likely to be used than one that is slower but is considered fun. Similarly, it seems obvious that the person's history and background is considered: it is unlikely that when presented with a new key based interaction someone who is proficient using a keyboard will react in the same way when presented with a new key-based interaction as someone with no previous experience with a keyboard. In fact, this principle explains why keyboard layouts have not changed significantly since they were designed, and laid out for, mechanical typewriters. Given these concepts, I am compelled to agree that the research upon which this thesis focuses should include the "Messy".

This inclusion leads to the question: how do we conduct research that includes such "Messy" elements? If we look towards traditional HCI it is difficult to find methods that can incorporate this concept of "Mess" since much of the relevant work is conducted through a positivist[1] lens which favours the removal of "Mess" in favour of repeatability. However, if we instead look to design, we find that methodologies that include the "Messy" are the norm for many researchers (Cross, 2006). It can be argued that the designerly approach is more suitable when confronted with these messy problems, as it is by nature an evolutionary process that evolves and redefines its goals and questions and is therefore inherently more flexible, and accommodating of change. Therefore, I am committed to situating this research methodology within a design paradigm.

There are many existing methodologies that encompass attributes that are important in this case, but I feel that no single existing research methodology would be a perfect fit for this work. As a result I have constructed a hybrid research methodology using aspects taken from a number of relevant methodologies. In the remainder of this chapter I introduce these relevant methodologies and discuss the aspects upon which I draw for my final methodology. I conclude with a discussion on the effect this methodology has on this research, its direction and findings.

2.2 Research *Through* Art and Design

Research within the context of design can be differentiated by what Frayling (1993) calls the "Big R" or "Little R" research. These terms essentially differentiate between that which would traditionally be called Academic Research (with an upper case R, typically used with reference to design) which focuses on communicable and peer re-viewable research, and research (with a lower case r, usually associated with arts practice) which, whilst shareable, is harder to communicate and peer review. These terms are further clarified by Frayling's later differentiation used to divide research in the area of art and design:

[1]discussed in-depth later in this chapter.

Research *into* Art and Design is "By far... the most straightforward, and . . . by far the most common". This includes historical research, aesthetic and perceptual research or research into a variety of theoretical perspectives, such as social, economical, political, material, and structural

Research *through* Art and Design is "The next largest category . . . is less straightforward than Research *into* Design, but still identifiable and visible." Examples of this are development work (e.g re-using a piece of technology for something no one had previously considered and communicating the results), Action Research (where a diary is used to explain a practical experiment in a step-by-step way, and a report is used to discuss the results).

Research *for* Art and Design is research in which the product of the research is an artefact (e.g a painting, where the thinking is embedded within). Whilst Research *through* Art and Design may also produce an Artefact, it is in addition to documentation (e.g a diary and report) and not a stand alone product.

Using this trio of definitions we can consider that "little r" research fits squarely within the Research *for* Art and Design, and "big R" Research covers both Research *into* and *through* Art and Design. Within this framework we are looking to produce academic knowledge: knowledge that can be peer reviewed and is easily communicable. This narrows our scope to both Research *through* and *into* Art and Design. Of these two categories, Research *through* Art and Design allows the closest fit with the previously discussed views on "Mess".

It is important to note that since Frayling's initial coining of the phrasings that make up this trio, the terms have been redefined by Findeli (1998) (Research into/for/by Design),Jonas (2006, 2007) (Research about/for/through Design) and Zimmerman et al. (2007). Whilst these alternative definitions may add clarification in some cases, they are largely non-incremental with respect to Research *through* Design and the description that Frayling offers in relevance to this thesis. What these re-definitions focus on is the removal of the "little r" - which Frayling refers to as Research *for* Art and Design - or the appropriation of these terms for a more positivist approach. For the sake of clarity, I am referring to Frayling's 1993 description of Research *into/for/through* Design when I use these terms from this point onwards.

As previously discussed, Research *through* Design involves the design and creation of an artefact or artefacts in combination with documentation (e.g. a diary, report, or published works). Agnew (1993) says that Research *through* Art and Design (of products) is often:

"hindered by the lack of any fundamental documentation of the design process which produced them. Too often, at best, the only evidence is the object itself, and even that evidence is surprisingly ephemeral." (Agnew, 1993, p. 1)

It is clear that for productive research using the Research *through* Art and Design methodology, documentation must be considered a necessity - an idea that is shared by many academics in the area, such as Gaver and Bowers (2012) when they discuss "Annotated portfolios". Without this documentation we risk the move from Research *through* Art and Design into Research *for* Art and Design - a transition from academic 'big R' Research into non academic 'little r' research. With reference to this thesis, Research *through* Design promotes the use of artefacts as tools for research, but only if the construction of, and interactions with, these artefacts are well documented. As a result, my methodology can involve the production of artefacts, but it must prioritise accompanying documentation.

Although Research *through* Design has been discussed here, it is subject to individual interpretation - each researcher or designer has their own interpretation of what Research *through* Design means to them, and their way of working. This is reflected by the definitions and redefinitions that these terms have received over the years, e.g. the disparity between Zimmerman et al. (2007) (more focused in computer science) and Frayling (1993) (arts and design) for example. Similar, but different uses of the term can be seen when comparing and contrasting existing works that all reference Research *through* Design as a methodology.

This tension is apparent discussed in relevant literature (Forlizzi et al., 2011, Stolterman, 2008, Zimmerman and Forlizzi, 2008, Zimmerman et al., 2010) which tend to criticise the popular interpretation of the technique as being "unscientific" due to the embodiment of designers judgment throughout the process with suggestions that it would benefit from actionable metrics upon which to measure success. In contrast Gaver (2012) suggests that this is perhaps impossible:

"[traditional sciences] need a shared paradigm to make progress, or already have a shared paradigm and need to recognise controversy as a sign of progress . . . the reason that research through design is not convergent is that it is a generative discipline, able to create multiple new worlds rather than describing a single existing one. Its practitioners may share many assumptions about how to pursue it, but equally, they may build as many incompatible worlds as they wish to live in"(Gaver, 2012, p. 943)

In other words, the generative nature of process of design itself encourages different and perhaps incompatible interpretations. I position myself similarly to Gaver's understanding of the phrase (Gaver, 2012). As such, it is appropriate to clearly outline what my personal interpretation for the term is, complete with epistemological foundation, relevant methodologies and what factors I aim to control to achieve research using this methodology.

The remainder of this Chapter discusses this in further detail.

2.3 Constructionism & Constructivism - An Epistemology

One of the primary difficulties with Research *through* Design is that the experience and subjectivity of the designer can often take a leading role, and the process and result are affected by the culture and the tacit knowledge held by the designer. Therefore, how do we present a case that knowledge generated throughout the design process is valid knowledge? To answer this we must consider our epistemological standing. If we take a positivist approach to this question, it is likely that it will be impossible to answer. Positivist approaches consider only knowledge gained through empirical evidence or mathematical reasoning and logic as valid. However, according to Crotty (1998, p. 47-48), it is impossible to separate consciousness and subjective thought. In this case, any research is going to occur though the subjective lens of the researcher (regardless of intent) and is therefore invalid knowledge according to a truly positivist approach. An alternative view of the world that supports a socially constructed reality is Constructivism, a term first coined by Piaget (1950) (this work was not translated into english but see work of Piaget and Duckworth (1970) and Piaget and Wells (1972)). This view is encapsulated well by Balbi (2008) who describes Constructivism as:

"an epistemological premise grounded on the assertion that, in the act of knowing, it is the human mind that actively gives meaning and order to that reality to which it is responding"(Balbi, 2008, p. 16).

This argues that there is no 'valid' or 'true' interpretation of the world, but instead only a subjective interpretation held by the person experiencing it. On the same premise, Constructivism argues that no object has meaning ascribed to it, the meaning originates from the interaction that takes place between the object and the person. HoweverRamirez (2009) argues that Constructionism, rather than Constructivism, may be a more appropriate fit for Research *through* Design. Rodriguez states that Constructionism:

" calls for a balance between the information we can gather from the object in the world, and the interpretation we construct from this information. However, if I use my designing as a method of research and the insights I gain from it as part of my data ... the information I am gathering ... is more subjective than objective"(Ramirez, 2009, p. 6)

Constructivism whilst similarly built upon individual interpretations of the world, focuses on the "meaning-making activity of the individual mind"(Schwandt, 1994, p. 127); that is, every individual's view of the world is as valid as any other, including the designer or researcher themselves.

With reference to the research area and topic discussed in Chapter 1, looking at this research through a Constructionist lens means that we can openly incorporate and acknowledge the (possibly tacit) knowledge of the designers and the researchers throughout this research process. This means that, as the researcher, I am able to contribute during the research process without fear of biasing results. However, I do have to be clear about my motives and biases: it is not enough for the research to comprise just work based upon this tacit knowledge from the researchers or designers, it must be part of an overall picture that is a balance of this knowledge and the knowledge gained from external sources. As a result I must make it clear upon which knowledge each decision is based, and explain why in detail.

2.4 Postmodernism - A Theoretical Perspective

Postmodernism is a reaction to modernism that brought with it a questioning of the previous approaches to knowing. Whereas modernist approaches generally rely on a single method of knowing, such as empirical evidence, post-modernism advocates epistemological pluralism, which inherently supports multiple ways of knowing. Postmodernism suggests that rather than one cohesive reality - which can be measured and adheres to certain laws - reality exists as a number of narratives fragmenting "into disorderly array of little, local stories and struggles, with their own, irreconcilable truths" Maclure (2002, p. 62). In the case of design, the postmodernist argument would be one that values the detailed description of truths and stories happening within these fragments as above or equal to quantitative measurement. Oakeshott (2015) suggests that "it seems self-evident that practical engagement in an activity should be a prerequisite to having a knowledge of that activity" (further discussed by Jackson (1996)), explaining that experience with an activity is at least part of generating related knowledge. This view aligns well with Research *through* Design, - it prioritises the construction of knowledge through practical engagement and the activity of design. If we combine the idea of fragmented realities with the concept of epistemological pluralism, it is logical to conclude that postmodernism supports tailoring research methods towards the reality fragment in which the research is being conducted. With reference to the research area and topic discussed in Chapter 1, this means that we can treat each aspect of research as existing within its own "fractured reality" rather than being part of a global set of truths in a cohesive reality. Therefore, we can tailor the research method to suit the specific reality that we are investigating. It is also important to provide some context in order to define the reality that we are investigating.

2.5 Action Research - A Reflective Practice

Action Research is similar in approach to postmodernism as it acknowledges the non-linear relationship that is often found between cause and effect, and the existence of complex social phenomena that are difficult to explain. Rather than maintain a barrier between the researcher and the object of research, Action Research aims to "abandon the notion of understanding as a product of the enterprise of a lone researcher, and to engage stakeholders into the research process"(Bradbury-Huang et al., 2009, p. 123). This approach seems particularly valid for Design *through* Research, Constructivism and Postmodernism as Action Research allows us to engage with stakeholders in order to gain a deeper understanding of their realities. It also acknowledges the value of the insights that the researcher, or designer brings.

Additionally, Action Research often exists as cyclical process. One of the most influential models is that of Kemmis and McTaggart (Denzin and Lincoln, 1994, Kemmis et al., 2013, McTaggart, 1989, 1991) that defines a process of plan, act, observe, reflect that then, after gaining more knowledge and insight, brings the process back to the planning stage. The value of this more reflective approach is also discussed by Schön (1987) in his book, Educating The

Reflective Practitioner, which explains how professionals meet the challenges of their work with a kind of improvisation that is improved through practice. This allows knowledge to be incremented at each stage of the process, and the research or design to be adapted in order to best accommodate any trajectory changes. It is possible to draw upon Action Research and emulate the iterative research cycle in this research methodology. This method allows the most flexibility to adapt to direction change and new parameters for the research. Similarly, other aspects of Action Research mirror aspects of the design process, such as engaging with stakeholders and acknowledging the insights provided by the designers and researchers. Appropriated, these aspects help to construct a methodology which better aligns with the Postmodern perspective I have outlined in the previous sections.

2.6 Grounded Theory

Grounded Theory with a Postmodern turn can be a useful methodology for assessing the lessons learnt from Research *through* Design (Ramirez, 2009). There are a number of interpretations of Grounded Theory sparked by the methodological split between Glaser and Strauss (2009) following their original publication which introduced Grounded Theory. These have been discussed elsewhere (Ralph et al., 2015). This discussion initially aligns with the "Strassian" view on Grounded Theory, which is already positioned with a more Postmodernist vision when compared with Glaser's. Grounded Theory suggests that rather than performing a literature review, producing a hypothesis, and performing study to prove or disprove this hypothesis, research may begin instead with data collection. This data is then inductively analysed to create new theory through a process of theme identification, coding, and grouping (Clarke, 2003, Corbin and Strauss, 2008, 1990, Strauss et al., 1990). As a result, this is most useful when the area of research is non-specific. Corbin and Strauss (2008) say that it is suitable when:

"all of the concepts pertaining to a given phenomenon have not been identified, or arent fully developed, or are poorly understood and further exploration on a topic is necessary to increase understanding" (Corbin and Strauss, 2008, p. 25)

The concept that we can research an area without being able to completely define a question is well aligned with the methods and processes of design. However, traditional views of Grounded Theory and even the "Straussian" view of Grounded Theory adhere to a reductionist explanation of reality (Ramirez, 2009). Clarke (2003) builds upon this view of Grounded Theory giving a further push towards Postmodernism. In Clarke's vision of Grounded Theory the researcher ("we") acts as a "mapmaker" working together with the stakeholders or participants ("they"). Clarke (2003, p. 31) says that "perhaps the radical reflexive act we perform as mapmakers is to reveal ourselves in and through analysing what 'we' do as well as what 'they' do". In Clarke's vision both the "we" and the "they" components are crucial to the construction of knowledge, and Clarke suggests the use of "Situational Analysis" maps to try to record the relationship between all stakeholders involved. Clarke remarks that this map is not meant to overcome the messiness, or necessarily be coded and analysed as in traditional Grounded Theory, but is instead used for practical reasons - "Some people may not even want to do the ordered working version. That's fine. It isn't necessary" (Clarke, 2003, p. 89). This "Clarkian" view aligns more closely with Research *through* Design as it suggests a more Constructivist view of knowledge: meaning is constructed, rather than truths discovered, and multiple fragmented stories or points of view (researchers, individual stories from participants) are used. What we can take from the "Clarkian" view of Grounded Theory is a focus on an inductive analysis process with the goal of discovering emergent themes. This approach implicitly supports the research of a research area rather than a specific question; we cannot predict ahead of time what themes will emerge throughout the research.

It is clear that "Grounded Theory" experiences similar definition fluidity to that which accompanies Research through Design: multiple definitions and interpretations exist depending on the backgrounds and epistemological positioning of those interpreting it. Whilst the "Clarkian" view - one which is more post-modern - is more fitting given the described epistemological positioning, even this is not completely appropriate for this work. Although many consider Grounded Theory as a qualitative research methodology, it should instead be considered as an inductive methodology. As such whilst this work is not explicitly adopting a grounded theory approach I am acknowledging the parallel with research through design. It is the concept that context specific research is appropriate rather than larger global narratives that is assimilated for this work rather than the (often context specific) methods of textual data coding and analysis often associated with Grounded Theory. Thus it could be said that this research adopts a grounded approach rather than performing a grounded theory.

2.7 Qualitative & Quantitative Data

At this point it is worth considering the types of data that are appropriate to these methods. A Postmodernist approach supports epistemological pluralism, meaning that data collected for the research can be acquired using a variety of methods, and be qualitative or quantitative in nature. Whilst there will be instances that the specific reality that we are investigating may require the collection of quantitative data, within the Postmodernist Constructivist framework outlined above it makes more sense to focus on the qualitative. It is the goal of this research to construct knowledge through an iterative (and reflective) process of design, which focuses on the stories, and emergent behaviour, of people using the artefact or system being designed. As such at the start of this process we cannot foresee the path of research and therefore it makes little sense to assume something about the data to be collected as this would enforce some element of bias to the research path. Quantitative data collection requires the researcher to make such assumptions about the direction of research in order to collect data. This is often impractical for research focusing on emergent topics as it is not known ahead of time the direction of the research, and, accordingly, the metrics to be measured. Given this discussion, qualitative information is likely to be the most prominent source of data, however the research process will consider other forms of data where appropriate.

2.8 Summary and Conclusions

Throughout this chapter I have discussed the view of Law (2007) on “Messy worlds”, and how agreement with this view means that many research questions are reduced to much more ambiguous research areas, similar in nature to Rittel and Webber’s “Wicked Problems”. It was then discussed how designers are often faced with this type of problem, and how the design process deals with this. From here, Design Research was divided into a number of categories, Research (big R), research (little R), Research *into* Design,Research *for* Design and Research *through* Design. From here relevant methodologies, epistemology and theoretical perspectives are explored that exist within the literature to support a Research *through* Design approach. From this, it can be concluded that Research *through* Design can be an effective research tool from within a Postmodernist/Constructivist framework if it meets the following criteria (based upon the work of Ramirez (2009)):

a) Well documented

Research within the arts and design field has potential to exist entirely as an artifact that is created during the design process. All knowledge that is gained throughout this process is embodied within this artefact. Law and Angrew have both discussed the difficulties of a "little r" research process, and state that the way to move into the academic "big R" Research is to effectively document the process. This view is shared by Gaver and Bowers (2012) when they discuss Annotated Portfolios. This documentation should include motives and describe how each decision in the process was reached.

b) Connected to the outside world (be more than my own experience).

The Constructivist approach acknowledges that the designer may produce valid knowledge, thought it is usually deemed as subjective or tacit. However, the Constructivist approach also suggests that the subjective knowledge sourced from the designer or researcher is only valid if it is used in combination with knowledge from the outside world i.e. participants or users. This view is supported by Clarke's view on Grounded Theory when discussing the research in terms of "we" and "they".

c) Subject to an analysis of any designs realised

As we are working with the goal of producing "big R" Research, the research process does not conclude with the creation of an artefact (as with "little r" research). Any artefacts produced are subject to a documented analysis after their creation. This step is almost inherent as we are drawing upon aspects of the Action Research cycle; the artefact will need to be analysed (the "reflect" stage) in order to begin the next cycle: "plan". This is also supported by the Grounded Theory approach that promotes the generation of knowledge from an inductive evaluation after data has been collected.

Furthermore, by extracting elements from a number of discussed methods and perspectives, a "hybrid" research methodology can be built that best suits the research area discussed in this thesis, which is well grounded in relevant theory.

Specifically, an amalgamation of the points highlighted earlier in the chapter allows me to describe the methodology used in this work as: an incremental research cycle that exists as a well documented process of Artefact creation, analysis and inductive reflection. The research conducted throughout this process is focused on the investigation of the emergent stories and truths of the individual's experience through interaction with artefacts. The research data is a balance of data collected through contact with the outside world and the tacit knowledge and insights contributed by the designer. As this process is focusing on emergent themes, it is difficult to be specific about a singular research question; instead a research area is defined as the starting point. Additionally, this process will involve predominantly qualitative data collection han it will quantitative. Figure 2.1 shows this process in its entirety.

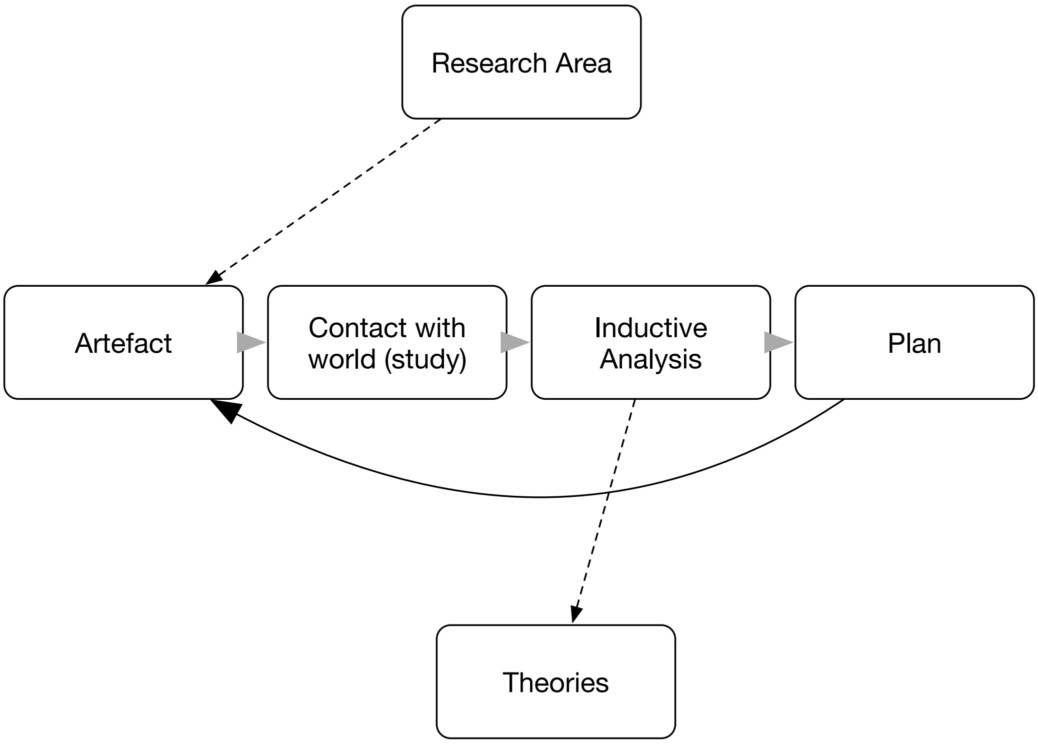

Figure 2.1: A process diagram for this hybrid methodology.

As discussed in Section 2.4, this process advocates multiple ways of knowing. This concept is supported by Mixed Methodology (Creswell et al., 2003) which promotes using different data types, methods, methodologies and/or paradigms during studies, and indeed Research as Bricoleur (Gray and Malins, 2016) which supports the use of any data sources available to draw conclusions. What this means for the research process in this thesis is that each iteration through the cycle will include a method tailored to that research. Therefore, in addition to the overall (macro-level) methodology, process and literature reviews given in these earlier chapters of the thesis, each cycle (represented as a separate project) will also include an additional (micro level) methodology, process, and literature review where necessary.

It is also worth noting that although throughout this chapter the reasoning and choices behind the creation of a hybrid methodology for doing research have been discussed, aside from stating that this process must be well documented, up to this point very little focus has been given to the artefact creation process itself. Verplank (2009) describes the basic design process as cyclical, and comprising distinct stages as seen in Figure 2.2. At the core of this process is "craft", which is a cycle of "hunch" and "hack". Ideas allow the generation of alternatives that can bring us to prototyping,testing and eventually to a product (artefact). Verplank states that "Design is the 'transfer orbit' that gets us out of a small orbit into a larger one", referring to the ability for the designer to generate ideas, moving from "craft" and into design allowing the production, testing and evaluation of a number of alternative prototypes. As such in this work the artefact design and implementation process will involve some level of the "craft" cycle ("hack" and "hunch"), but it will also contain idea generation that leads to consideration of alternatives, and justification of choices throughout the process of prototyping,testing and the eventual creation of the artefact.

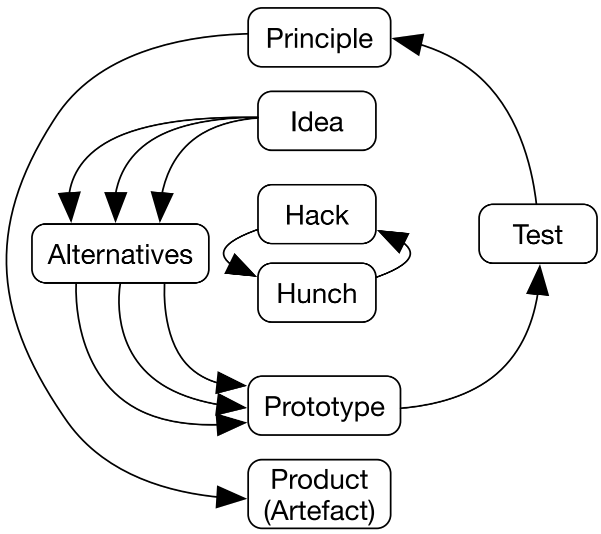

Figure 2.2: A portion of Verplank's product design process diagram.

This research process dictates that I start with a research area, in this case emergent interaction behaviour in physical/digital spaces. From here I design and construct some artefact to be used as a prop/talking point when interacting with users and the outside world. Throughout this interaction I perform data collection that I later inductively analyse for themes and concepts that can inform the next iteration of artefact design. After a number of iterations the research moves out of the cyclical phase and turns to an inductive analysis of the data collected as a whole, presenting this information as theories and frameworks to the academic community. It is important to note that through this process the artefacts designed and built are important to the research, but the research is not focused upon them. The artefacts should be considered a means to an end, or a prop to permit investigation.

Additionally, it is important to note that the part of this process labelled “Contact with world (study)” involves some sort of interaction between the outside world and the artefacts. However, as discussed in the previous sections, the path of the research throughout this process is dynamic and reactive to the data as it is collected and problems are refined. It is therefore difficult to know in advance the specific format of this contact with the outside world as it will be tailored to the artefact and truths being investigated. However, it can be expected to exist as one, or a combination, of the following formats:

Workshops

People are invited to interact with the artefact as a group or as an individual and then join a discussion about the artefact, their experience interacting with the artefact itself, about and any insights they may have. The background of workshop participants is also discussed in order to give their views context.

(Informal) Interviews

Much like a workshop setting, but purely on a one to one basis. People are invited to take part in a discussion about their experience interacting with the object.

Interventions

In the case of interventions, a person’s normal work flow for a particular problem is disrupted in some way. The research then focuses on the person’s reaction, experiences of, and new solutions to the problem.

In conclusion, the aim of this chapter is to clearly outline how, and why, this research is carried out the way it has been. When facing a problem that can be classified as “Human Computer Interaction”, it is easy to see that much of the relevant work exists within a reductionist view of the world. More recently however, a more Post-Positivist approach has been gaining momentum within the research community, and this view is one which with which I align myself as a researcher. This process is one that more closely mirrors the design process, and in this chapter I have formalised this view as “big R” Research that fits within Frayling’s Research *through* Design category This view supports Law’s idea of “Mess”. Furthermore, I have gone on to discuss how the Constructivist epistemology and Post-Modernist perspective are supportive of Research *through* Design, and additionally I have described how an amalgamation of these perspectives, combined with aspects of Grounded Theory and Action Research can produce "valid" research in this area. I have then formalised what this hybrid methodology means, in terms of specific goals and guidelines, and in the format of an overall process diagram that will be followed and referred to in the rest of this thesis.

Chapter 3

Research Project 1:DTF

3.1 Introduction

As previously discussed, this work is presented chronologically with a number of smaller, incremental projects which aim to explore the research area with methods appropriate to the Postmodern Constructivist methodology discussed in Chapter 2. Whilst interesting in their own right, these projects all contribute to the overall understanding of this hybrid Phygital space. The next three chapters act as an annotated portfolio comprising three projects which have been presented as papers for academic consumption. These chapters expand upon the work presented in the papers - links to which are included in the appendices of this thesis and will be available to download for the foreseeable future. In these chapters, the projects are explored with more thorough representations of the work, and providing additional details and analysis where appropriate - especially regarding aspects of the research that *didn't* work out as was initially intended, or the hack & hunch process. Each project is then discussed in terms of its themes, and what can be taken into the next project (or the next iteration in the process cycle). A more thorough, all-inclusive inductive analysis of this work will be discussed in Chapter 6. In order to signpost the progression through this process, diagrams such as those seen in Figure 3.1 are positioned throughout each chapter.

3.2 Planning

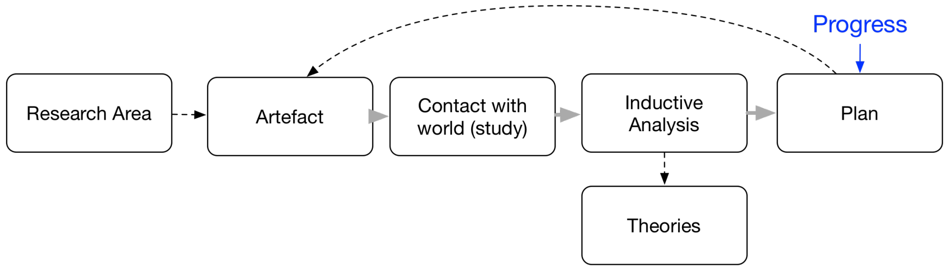

Figure 3.1: Current progress in the design cycle.

The Digital Terra Firma (DTF) project produced the first artefact of importance to this research. At this point in the research cycle (see Figure 2.1), the only information about the research area that was available was that of the problem area : "interactions in hybrid physical/digital spaces". Therefore, the artefact needed to be as flexible as possible in terms of allowing emergent behaviour that would then direct the rest of the research. This artefact needed to exist in both the physical and digital realms, allow open-ended interaction with this space, and both physical and digital objects within it. In this section, the artefact creation process of the DTF system is discussed (as outlined in Figure 2.2), and choices used through this process are described and explained.

In order to begin designing and building this system the relevant goals and limitations had to be considered. In order to achieve the goal of synchronising a space physically and digitally, a system needs to be able to somehow detect the physical properties of a space, whilst simultaneously augmenting it with the virtual. Figure 3.2 shows the flow of information through the system which was initially proposed.

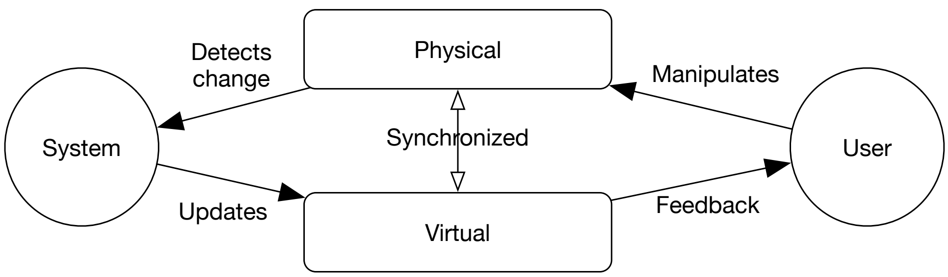

Figure 3.2: The synchronised system data flow.

A preliminary 'brainstorming' session was held to explore potential ways that this synchronisation could be achieved. This session took place in a meeting room inside the LICA (The Lancaster Institute of the Contemporary Arts). The building acts as a base of operations for multiple cross disciplinary research projects and PhD students. For the session, people within the building were invited to attend. As a result, those that attended had a wide range of experiences and backgrounds, although they could typically be described as having a focus on making and experience design.

In general ideas were divided into two distinct categories: those which discussed ways in which the virtual could be presented to the user (i.e physical synchronised to the virtual), and those which discussed ways in which the physical could be detected or tracked as to be useful in the virtual (i.e virtual synchronised to the physical). A few of these ideas can be seen in 3.3. Although many of these ideas were trialled throughout the "Hack and Hunch" cycle within artefact design, success was achieved to varying degrees. In many cases cost or reliability were prohibitive to the continuation of the technique.

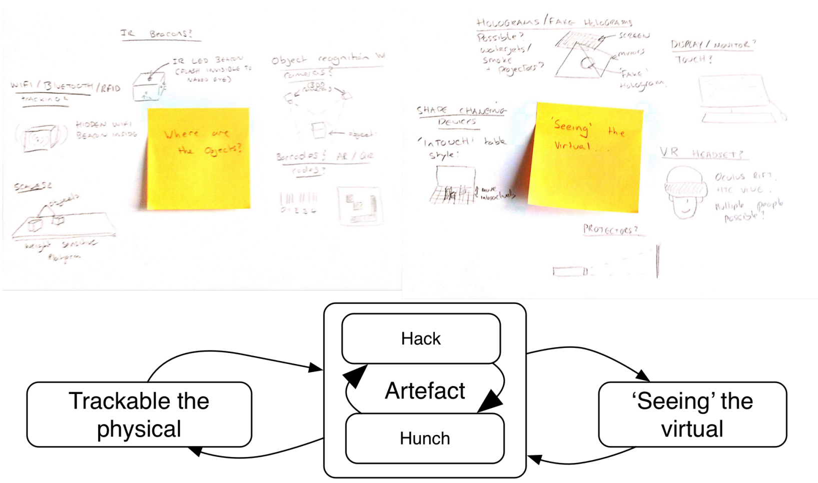

Figure 3.3: A few of the ideas from the workshop separated by category.

The first consideration is the way in which feedback is presented to the user. In an ideal world, a technology similar to a 'hologram' would be used - a common technology used in science-fiction in which digital entities can be made to exist in the physical world, with varying degrees of physical presence.

However, as no such suitable technology is feasible for this work, alternative approaches must be considered. For this I highlighted three main candidates: VR goggles, a computer screen, and projectors. All three of these are readily available at the time of writing, and each offer some way of supplying digital visual information to people in different ways. The use of VR goggles was quickly discounted as, without significant effort, it limits interactions to a single user and the hardware is typically much more expensive that the other available options. As the goal here is to look at emergent behaviour, the possibility that multiple simultaneous users may be a factor should be considered. Either a screen or a projector allow simultaneous use by multiple viewers. Whilst the screen, as a physical device, is potentially more familiar to users, the projector is more flexible as it can display information on any surface: it is not limited to smooth flat surfaces as with a screen. Despite this, it was decided that both a screen and a projector would be used so as to not limit the scope of the work. The projector would be used to augment physical objects with digital information. Any information that was infeasible to display via the projector, or information that was not tied to a specific location within the space would be portrayed to the user via the screen.

The next design aspect to be considered was the way in which the system would sense and detect changes within the physical world. For this many options were considered - the following list is a subset of the most applicable: a standard Universal Serial Bus (USB) web camera; a depth camera, such as the Microsoft Kinect (see Figure 3.6), or Creative Senz3D (see Figure 3.5a); a leap motion controller (see Figure 3.5b); various shape change detecting media (e.g a mesh of bend sensors); and Computer Vision (with markers of some kind). A “hack” and “hunch” style cycle (see Figure 2.2) was used to quickly prototype and test the best options for this purpose. it was found that techniques for object detection based solely on Computer Vision (e.g webcam and object detection) were unreliable: changes in ambient room lighting and moving shadows often rendered any detection techniques inoperable. The leap motion was built primarily for the detection of hands, which it performs very well. However, the detection of anything else (e.g an object or surface) proved more difficult and less accurate. The concept of shape-change detecting media is an interesting one which may have worked for this premise. However, in reality the hardware that could be obtained seemed to be in its infancy, with no options readily available to use. Some tangible building blocks similar to Smart Programming Blocks (Kortuem et al., 2010) were considered, but it was decided that interacting with this kind of discrete media may limit interaction possibilities.

3.3 Artefact Creation

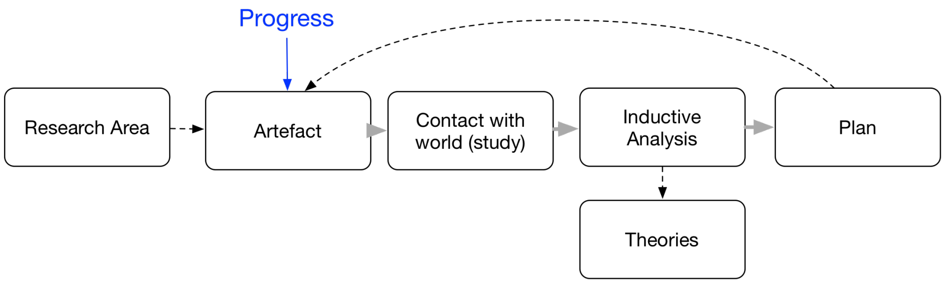

Figure 3.4: Current progress in the design cycle.

(a) Creative Zen3d.

(b) Leap Motion.

Figure 3.5: The leap motion and Creative Zen3D sensors.

Throughout this hack & hunch style design process, many alternatives were trialed as possibilities for going forward. It was concluded that the best way to record the digital state of the space was with the use of the Microsoft Kinect (Zhang, 2012): it is cheap, readily available with lots of programming support, and offers multiple sensors with which the state of the physical could be recorded. Whilst technologically similar to the Creative sensor, the Microsoft Kinect boasts a superior development community and programming interface.

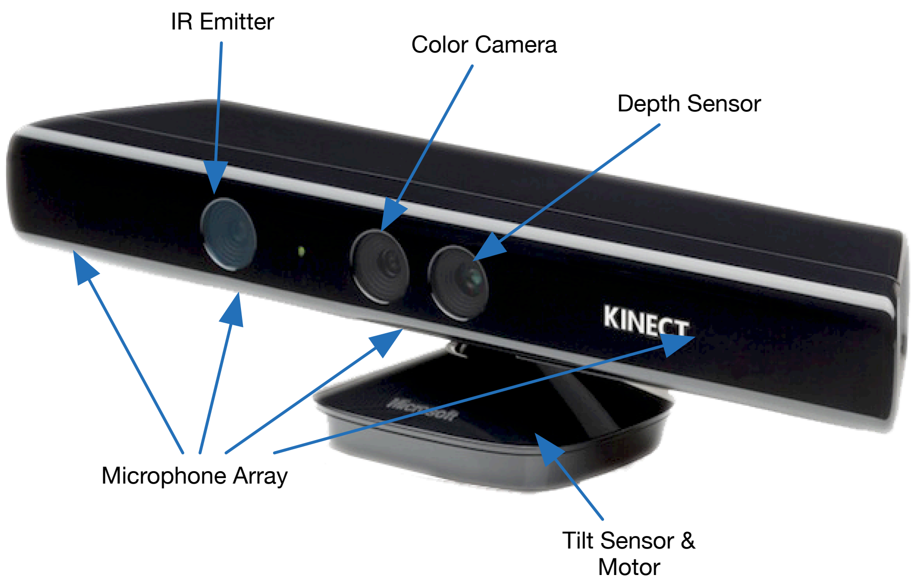

Figure 3.6: An illustration of the various sensors available with the Microsoft Kinect.

Now that some limits to the physical design of the system had been recognised (it must include a projector, screen, a Kinect, and some physical space) it was necessary to devise some context for this interaction, and put together an artefact that incorporated all of these elements. After discussing and exploring different contexts within with the interaction could occur, and exploring other projects that used similar technologies (Hardy and Alexander, 2012, Wilson et al., 2012), it was decided that the context that offered the most flexibility for this interaction space was one of virtual world design: a user can manipulate physical objects in the interaction space and a virtual version will be synchronised with this physical space. Information about the virtual space would then be projected into the physical elements that the user is manipulates, and an overall space would be shown on the screen, situated next to the interaction space. Example use cases for such a system are video game level design, interior design, and landscape design amongst others.

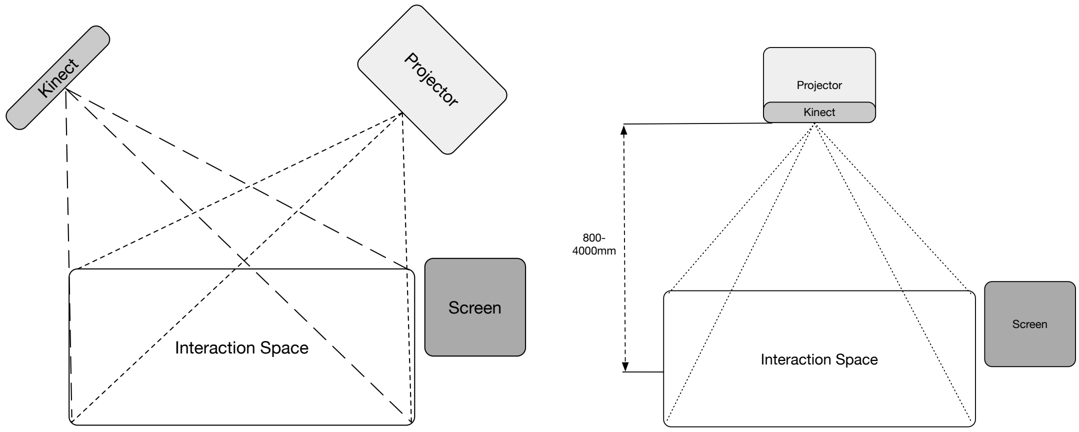

(a) Initial interaction space concept.

(b) Actual layout after testing

Figure 3.7: An illustration of projector and sensor positioning.

Once an initial concept had been chosen for how this space would work, a frame to hold the projector and Kinect in position above a table that would be used as the interaction space could be constructed. Although this frame went through a number of small iterations with minor improvements in terms of stability and size, the general layout of the DTF system is reflected in Figure 3.18a. One of the notable aspects that was tested and modified through the build process was the positioning of the Kinect and projector above the work space. Whilst the projector can theoretically be positioned anywhere above the space and adjusted using Keystoning[1], initial testing showed that the projector worked best positioned directly above the work space as this minimised shadows on an uneven surface.

[1]Keystone correction: a function that allows multimedia projectors to skew the output image, thereby making it rectangular on non perpendicular surfaces.

(a) Initial projectors used. **(b)** Epson short throw projector

Figure 3.8: The projectors trialled throughout this process.

Similarly, if the projector was positioned off to one side, the projected images would tend to beam onto the side of objects rather than onto the top, which proved confusing for users. Similarly, the Kinect could theoretically be positioned anywhere above the work space. However, the data produced by the Kinect was cleanest and easiest to analyse when positioned directly above the workspace. Additionally, due to internal hardware limitations of the depth sensor, the Kinect had to maintain a distance of between 800mm and 4000mm from the work surface[1][2]. The final layout of this system is shown in Figure 3.18b, whilst Figure 3.9 shows a picture showing one incarnation of the project set-up. Although projectors are all similar in their basic principles of their functionality, I found that short throw projectors were the most flexible as they did not have to be positioned so far from the surfaces that onto which they were projecting. Figure 3.8 shows two of the projectors trialed throughout the process, which were eventually surpassed by the very short throw projector model 'Epson EB-410We' shown in Figure 3.8b.

[1] https://support.xbox.com/en-US/xbox-360/accessories/sensor-placement

[2] forum discussion explaining this: https://social.msdn.microsoft.com/Forums/en-US/53a62e18-6acf-466b-9549-9692ff183d77/minimum-depth-distance?forum=kinectsdk

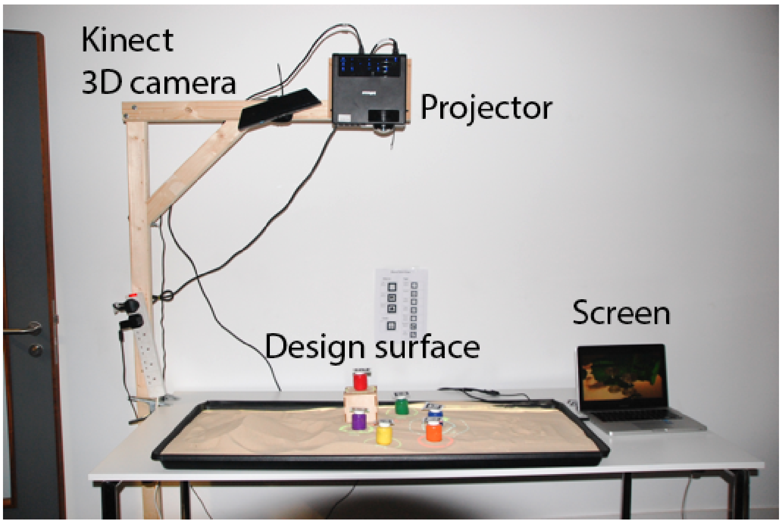

Figure 3.9: The overall setup for DTF, using a sandpit and coloured pots as design media.

3.3.1 Software

Once the physical system was set up, the software that would run the system could be chosen. The only hard limitation for this was the drivers available for the Kinect. As a Microsoft device, the only officially supported versions are available for the Windows platform and are written in C# [1]. However, at the time of writing there also exists an alternative unofficial driver called "libFreeNect"[2] that is multi-platform (works on most operating systems) and can be programmed in many languages. Due to personal preference in programming languages, and past experience with different operating systems, it was decided that libFreeNect drivers would be used in conjunction with a platform independent Java wrapper which affords cross-platform freedom.

As soon as data is received from the Kinect in software it is apparent that the data needs some form of sanitization. Whilst the Red, Green, Blue: a system for representing the colors to be used in an image (RGB) data images are usually free of noise and erroneous data, depth data gathered using the Kinect was prone to noisy and missing data caused by reflective surfaces, and various other differences in refractive properties of the area in view. To remedy this a small middle-layer program was constructed to interface between the libFreeNect drivers and an interface that the rest of the software written throughout this thesis supports. This middle

[1] https://msdn.microsoft.com/en-us/library/hh855347.aspx
[2] which can be found at https://github.com/OpenKinect/libfreenect

layer is shown in Figure 3.10. The object that does the clean-up is labelled as "Data Sanitiser". This object performs a number of Adaptive Directional Filters(Le et al., 2014) to sanitise the data, and replace any missing data within reasonable extrapolated values. This software module also "smooths" the data, by taking the average value for each pixel over an adjustable amount of time (this is known as a moving window average) and supplying these average frames to then be adaptively, directionally filtered, and passed back to the rest of the software.

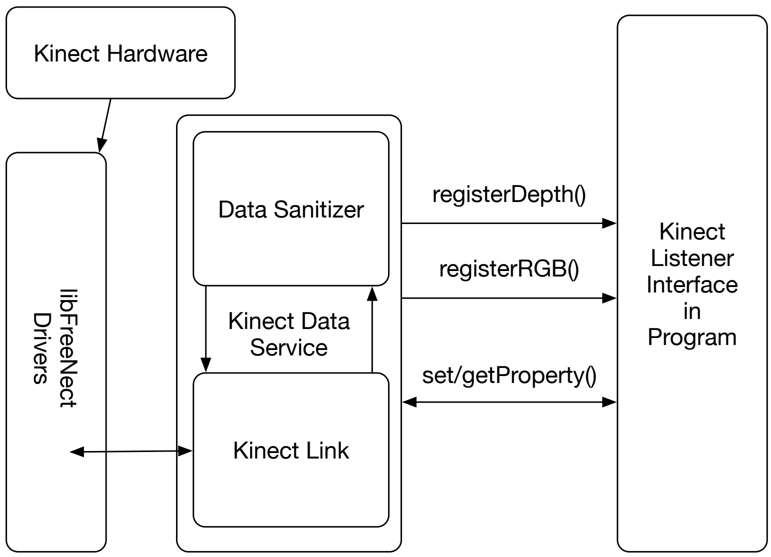

Figure 3.10: An diagram showing the data flow through software elements in the system.

This Kinect Listener Interface exists as part of a dedicated program that uses the Kinect inputs to create a virtual version of the physical interaction space in computer memory. After trialling a number of alternatives [1][2][3], the graphics engine JMonkeyEngine was chosen for this project. This was selected as it is written, and has an Application Programming Interface (API) for, Java in which the rest of the software is written. Additionally, it is very well supported in terms of online tutorials and an active community should any help be required. Most importantly it offers a good API for dealing with 3D models and model meshes. Since this had been used successfully previous projects of mine, I had experience and a library of existing code which could be used for this purpose. This means that once received, the cleaned data from the Kinect (both RGB and depth) could be used to construct a virtual model of the

[1]Unity: https://unity3d.com/
[2]UnrealEngine: https://www.unrealengine.com
[3]Cocos3D: http://cocos3d.org/

area - by linking both the colour and depth data. One of the first iterations of this process can be seen in Figure 3.11. In this, it is evident that the virtual representation contains the elements that exist in the physical realm: a keyboard (bottom right), a screen (top right), the stand (top centre) and a number of boxes on the left. The missing data smoothed away by the Data Sanitiser, and the effect of objects being too close to the camera can also be seen: the top of the stand (top left) which gets truncated.

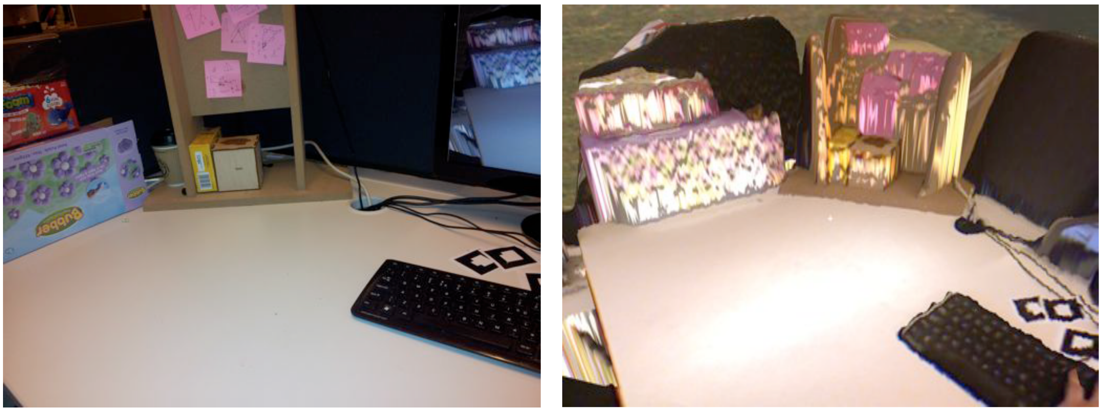

(a) Physical space. **(b)** Resulting Virtual representation of space.

Figure 3.11: A comparison of real a) and virtual space b) using the DTF system.

From here the investigation began into the ways users could interact with this system. Firstly, the scope for the manipulation of physical objects in the interaction space was considered. These manipulations are replicated in the virtual space. For example, a user can move, remove, or reshape various objects in the work space. They can even use their own body as part of the space: Figure 3.12 shows a user manipulating the virtual world by moving their arm in the physical world. In this, you can see that the user's arm is contributing to the physical shape of the mesh (shown in pink) and consequently their physical movements manipulate virtual objects (in this case, some virtual balls). This screen shot has been taken from one of the preliminary informal user tests that took place repeatedly throughout the designing and building of this system, and not all aspects were kept in the later versions of DTF. For example, the part of code that emulates physics and the laws of motion (gravity, collisions e.t.c) was later removed to reduce complexity and confusion for users.

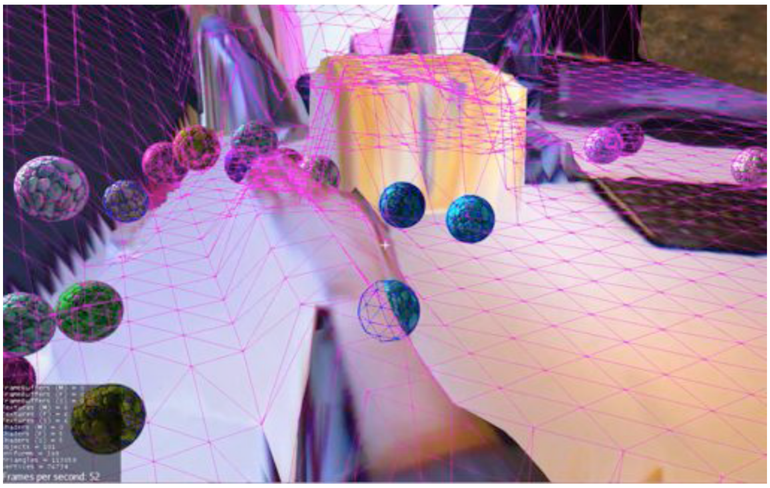

Figure 3.12: The virtual view of a user manipulating the space physically with their arm. In this case they are moving virtual objects around by knocking them with their arm.

Whilst this kind of interaction is useful in many ways, it doesn't allow much beyond the manipulation of the physical objects alongside a virtual representation of same. After preliminary testing it became clear that people did not always want a physical item to represent a tangible object. In some cases they wanted an object to represent itself, whilst in others it was expected that the object should be a tool of some kind. The work of Underkoffler and Ishii (1999) with the Urban Planning and Design (URP) toolkit suggests that object usages in this space can be considered as existing on a spectrum (see Figure 3.13): from an object as a pure object; to an object as an attribute; a noun (as in this case); a verb; or finally a reconfigurable tool (see Figure 3.13).

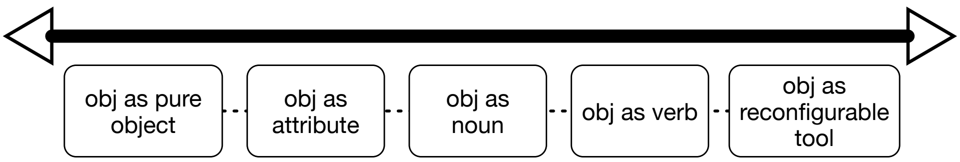

Figure 3.13: The object spectrum as described in URP.

This spectrum describes the way in which a user might want to represent an object in this space. On the far left *object as pure object* describes an object that is stripped of all attributes and descriptors including shape, colour and weight amongst others. The only attribute of importance is that it exists as an object. This may be useful for interactions in which the presence of an object is being detected, regardless of what kind. Further right on the spectrum is *object as attribute*, which describes an object stripped of all but one defining attribute, e.g colour, which may be useful for interactions that rely on detecting a certain attribute. In the centre of this spectrum is *object as a noun*: a literal representation of a physical object replicated in the virtual environment (the interaction type we have used thus far). Moving further right, *object as verb* describes the use of an object as a tool: something that does not manipulate a specific object but rather describes some other aspect. For example, in a virtual world designer, an object may manipulate a non-tangible element such as time or wind. Finally, at the far right of the spectrum *object as reconfigurable tool* exists: an object which exists as a tool that applies some effect to the space. This effect however can be reconfigured and is similar to the mouse pointer in a WIMP interface that performs different actions depending upon which menu bars and options have been activated. Although Underkoffler and Ishii (1999) is arguably using this spectrum to define explicit tools that are used together with the URP system described, this spectrum is effective at framing objects within this space.

These descriptions make it apparent that in order to offer the widest range of interaction possibilities, an object must be able to be "decoupled" from its physical properties and have the system detect in some way what the user is trying to represent with that object. For this, an exploration began into the concept of markers that can attached to objects in order to augment virtual behaviour, and that can be used independently to apply meaning without affecting the physical shape of the interaction space. Also investigated was the possibility of object detection using the camera and Kinect along side Computer Vision techniques (e.g OpenCV Bradski (2000)). Unfortunately, these methods of object detection proved to be very unreliable, especially with changing lighting conditions. Ideally these markers would be as unobtrusive as possible and as a result development turned towards objects that would communicate in some way that is imperceptible to users. This lead to experimentation with options such as Bluetooth, and other wireless protocols including infra-red touch detection through the table (Han, 2005). These approaches were more difficult than expected due to the need for additional sensors, and inaccuracies or interferences resulting from a research office environment (predominantly in the form of radio interference). Tangentially, the detection and positioning of wireless devices in such a small space has its own field of research (Patten et al., 2001), which is out of scope of this work. Similarly, the table top, whilst a nice interaction system, limited the amount of 3D manipulation that could occur on the surface. As a result the possibilities were narrowed down for these markers to those than can be visually detected by the system through one of the already available sensors on the Kinect i.e visually or infra-red (the Kinect depth sensor is based on an infra-red sensor). During the prototyping phase a number of markers were built that consisted of two infra-red LEDs and a simple circuit to pulse this LED with an encoded ID. The system would detect this ID and could then perform some action using the ID and the physical position of this marker. However, this approach was later abandoned in favour of AR markers due to simplicity of production, their smaller size and the fact that the Infra Red (IR) markers adversely affected the quality of the depth data received by the Kinect as they interfere with the infra-red sensor.

Figure 3.14: Examples of AR markers.

AR markers are much like 2D barcodes, or Quick Response (QR) codes but implemented specifically for the use within the AR context. Some examples of AR markers are shown in Figure 3.14. Whilst it is possible to use any image as a marker, for initial testing we used this barcode style of marker due to the more robust detection techniques for this design. These were generated using a mixture of softwares (Unknown, 2005c, 2014) to create both the marker, and corresponding .patt files. These .patt files are a description of the pattern, and are used by various softwares to detect the markers. Using a freely available AR marker library (NyArToolkit Unknown (2005b), based upon ARToolkit (Unknown, 2005a)) it was possible to incorporate the detection of these markers into the portion of the Data Sanitiser class that deals with RGB Kinect data. Figures 3.15 and 3.16 show some of the resulting detection when the AR markers were used. The different types of marker are discussed later in this section. It is worth noting that AR markers do not need to be a barcode-like image as any image will work to some degree, however, the best detection rates in research for this thesis were obtained using barcode-style markers such as those shown in Figure 3.14.

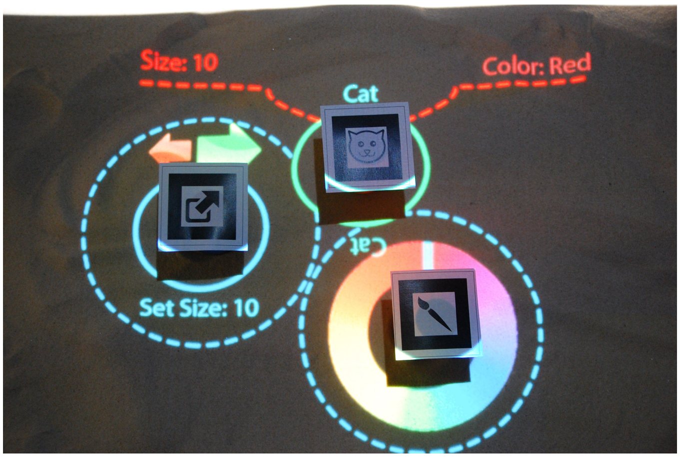

Figure 3.15: Example projections on system after tags have been identified.

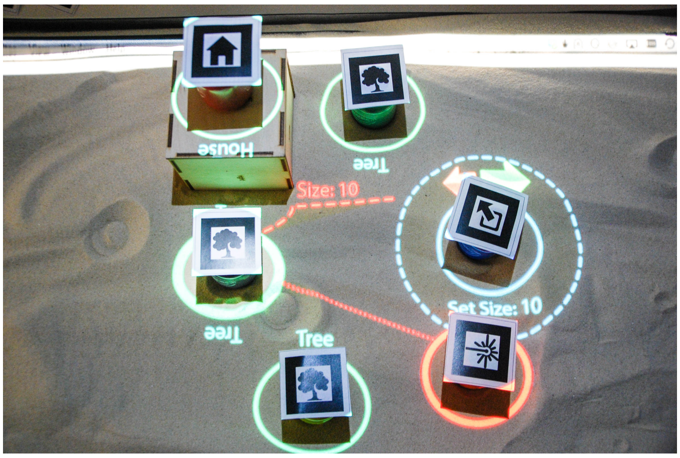

Figure 3.16: Identified tags, showing tags that represent objects (cat, tree) and tags representing tools (size, colour).

The inclusion of AR markers that are thin enough to be considered only 2D can now be used to embody a virtual element or attribute that has physical positioning but no physical body. Physical objects can also be augmented with some virtual element.

Within the realms of DTF and virtual world design, a number of basic interactions were implemented that users can perform when interacting with the system. These are shown in Figure 3.3.1.

Physical Object

A physical object is one that the user can move and manipulate in this physical space. Within the spectrum discussed above, this falls under the category of object as noun: it represents its literal self. This type of object can be used to construct the physical shape of the workspace.

Virtual Object

A virtual object is a marker object that can be placed somewhere on to the workspace by the user. As the marker is paper thin, it doesn't affect physical shape of the environment, but allows the system to detect its presence. This type of marker positions a virtual object into the virtual space though the marker resides in the physical. For example, placing a model house down in the virtual world. This is similar in nature to object as attribute on the previously discussed spectrum. These can be seen on Figures 3.15 and 3.15 as 'Cat', 'Tree' and 'House'.

Radial Parameter

A Radial Parameter is a marker that, when positioned, applies some sort of effect on nearby objects. When the object is removed, the effect is also removed. This can be considered an object as a verb. These can be seen on Figure 3.15 as the 'set size' marker on the left, and the 'colour select' tool on the right.

Directional Parameter

A directional parameter is a marker that acts somewhat like a proxy: when positioned, it applies the effect of nearby objects to those along a projected line. When the Directional Parameter is removed, the effect is also removed. An example of this can be seen on Figure 3.16 where the 'set size' effect is applied to a nearby tree. The Radial Parameter would not have been able to be positioned close enough to the tree without having an unwanted effect on other objects.

Paint Tool

A Paint Tool is one that can be used in the workspace to apply some attribute to areas within it by 'painting' onto it. Much like a physical paintbrush, these effects are permanent until overwritten. As the effect of this tool can be reconfigured by swapping out the AR marker (effectively changing the paint on the paintbrush) this object can be considered as a reconfigurable tool. This can be seen covering an area with grass in Figure 3.19 and further examples of this are discussed later in this chapter.

Meta Parameter

The Meta Parameter tool presence is detected anywhere on the workspace. if present it applies some global meta attribute to the space. For example, it might set time of day (and therefore sun position), weather, and temperature amongst others. The effects of one such marker setting the water level can be seen in Figure 3.21.

3.4 Contact with the World

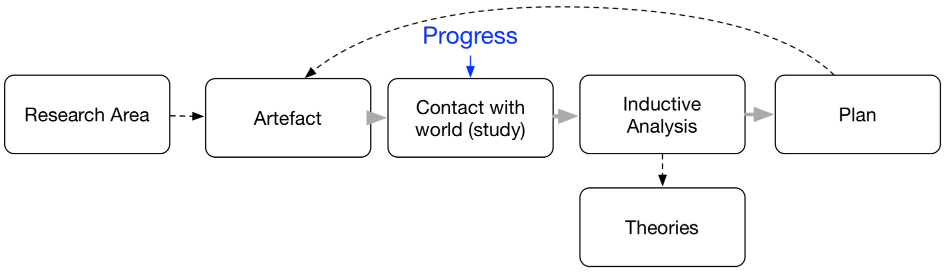

Figure 3.17: Current progress in the design cycle.

In order to begin to understand the way in which people would interact with this type of system, the DTF system was assembled in a public space within a well-populated and active research building on the campus of Lancaster University. The building is composed primarily of researchers, and university students, whose areas of specialism are extremely diverse due to the interdisciplinary nature of the research carried out in the building. As a result users of a wide range of ages, technical backgrounds, areas of expertise, and cultural backgrounds had access to the system. To begin, people were invited to interact with the system and familiarise themselves with the different ways of interacting with the space. This was explained as a sort of 'sandbox' style interaction where users can build something of their choosing, just as children can in a sandbox. Throughout this preliminary testing session, 15 users interacted with the system (six female, and nine male) of varying academic backgrounds and ages ranging from 23 to 50 years old. During the sandbox session, users were asked to narrate the process as they interacted with the system, following the Talk Aloud protocol (Lewis, 1982) (a method which asks participants to say whatever comes into their mind as they complete the task, thus highlighting reasons for choices). Additionally for this process, a 'toybox' was provided, which comprised various objects which could be used with the system. This toybox consisted of sand, Bubber[1] (A malleable clay link substance), PlayFoam[2] (Mouldable foam made of malleable

[1]http://www.relevantplay.com/Bubber/
[2]https://www.funlearning.co.uk/playfoam-original-modelling-beads

foam beads), toy tuilding blocks (various, including Lego bricks[1]), and other miscellaneous physical objects of various shapes and sizes. For the purpose of framing this work, the DTF system was posed as a a step in the landscape design process that may occur between user and expert (e.g a client and a landscape architect). Normally, in a landscape design scenario a non-expert user (the client) would explain their requirements to an expert (the architect), and the expert would interpret this information to produce some designs that are shown to the user at a later date: the whole process is very asynchronous in nature. The DTF system was framed as a tool which both the expert and non-expert users can use synchronously to streamline this process: a non-expert is able to physically model what it is that they want. Through the use of physical boundary objects (Star and Griesemer, 1989, p. 389) (an object which can be understood by all involved) communication between client and expert can be facilitated.

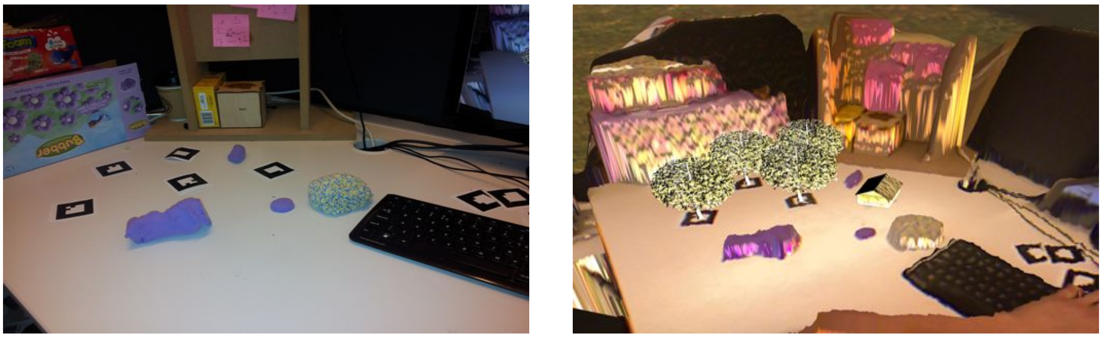

(a) Physical workspace

(b) Virtual Workspace

Figure 3.18: A comparison of the physical workspace, and the virtual workspace. In a) you can see the virtual markers, and their effects (model tree/house) are shown in b)

[1]https://www.lego.com/en-gb

(a) Physical workspace **(b)** Virtual Workspace

Figure 3.19: A comparison of the physical workspace, and the virtual workspace. In a) you can see the virtual markers and a radial marker, and their effects (model trees and grass) are shown in b)

Users were quickly competent at building and experimenting with physical building blocks as it "felt very natural" to them (as is to be expected given the heavy influence of TUI and NUI on this work). People would quickly build and design the shape of whatever it is that they were trying to build using anything they had to hand: items from the supplied tool box, or items that they had brought with them (drink bottles, a basketball, and in one instance a sandwich). They were also happy to work together with other users to discuss their design, and remarked that the physical nature of the interaction made this kind of cooperative work easy to undertake. Whilst this is encouraging, a number of users were also confused by the "extruded 2D" nature of the virtual landscape produced using the Microsoft Kinect. For example, when building a bridge (shown in Figure 3.19) the bridge constructed physically has a void beneath the arch, (3.19a) whilst its virtual representation has no void beneath it and is instead connected to the floor (3.19b). This is an artefact of the single point of view of the sensor (top down) and as a result it has no way to differentiate between a void and a solid object. On a related note, people were often a little confused when parts of their body were in the design space as it would get recreated virtually. For example, if a user is resting their elbow on the table, or reaching over to modify something in the work space, their arm will temporarily show virtually in the workspace as the system has no way to determine if an object is part of the user or not. However, after a little practice, users seemed happy with this process.

3.4.1 Technical Issues

Although not specifically related to the major outcomes of this research project, there were a number of technical issues that needed to be resolved throughout the hack & hunch process which arose from the system as described above. The majority of these were to do with the accurate detection of the AR markers. Whilst these markers are fairly easily detected with even lighting, once the markers are placed underneath the projector, the projected image falls on top of the marker which makes it harder for the camera to detect. Whilst a number of different approaches - such as alternative marker techniques (e.g infra-red beacons, infra-red ink e.t.c) - were trialled to resolve this, the most success was had by projecting a white area around the markers once they were detected, which simulated even lighting across the whole marker. In early versions this had difficulties if the detection was lost for one frame since the projected area would disappear, reducing the likelihood that it would be reacquired. To combat this, an area to the side of the system was designated as always white. This area was termed the "bench", and was where the markers could be placed if they were not being detected by the system properly. In later version of the system, the detection code was more robust so the bench was not needed.

3.5 Analysis & Themes

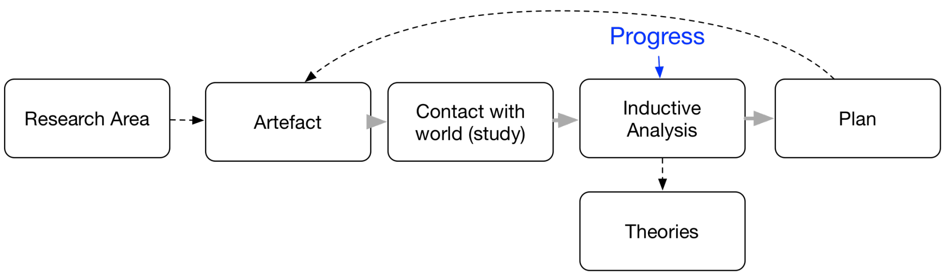

Figure 3.20: Current progress in the design cycle.

Through this section discussion has focused on the production of an artefact in the form of a tool for users to design a virtual landscape. Although the findings set out here are in reference to a specific artefact, the findings and emergent themes that can be taken from this research cycle in order to inform the next iteration have yet to be discussed.

Physical (Tangible) interactions are "easy" to learn and use

One thing that this testing and evaluation period has strongly supported is the notion that systems with a physical interface seem more natural to users. This is not to say that users don't require at least some learning to become proficient, but rather that a user is already comfortable with many of the actionable properties of an object before use due to their extensive interaction with physical objects in everyday life. This seems to support claims that the goals of this project are alongside those of TUI and NUI. Interestingly, some users in this workshop who described themselves as more of a "techno-phobe" remarked how surprising it was that they could learn to use the system so easily. Moreover, this physicality seems to make collaborating towards some common goal particularly easy: people happily manipulate physical objects to aid them in conveying solutions to problems to others with whom they are working. When asked, users said that this kind of physical manipulation comes naturally, whereas with a traditional GUI: "it is hard to do as you don't know how to use the program that they are using. Also, there is only one keyboard so you have to take turns if there are more of you.".

People expect a more multi-sensory approach

Throughout the process, it became apparent that although people were interacting with the system via physical manipulation and observing the results with vision, this was not everything that was expected by the users. This was initially apparent when one of the models that people could place was a flame object that would light up and cast shadows on the surrounding area (shown in Figure 3.21).

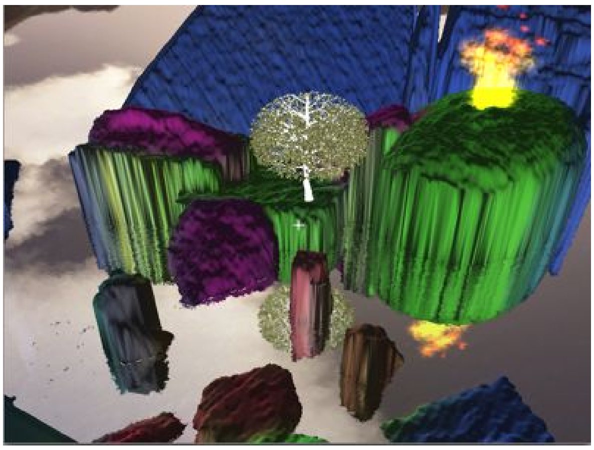

Figure 3.21: An image showing the placement of the 'flame' marker in a virtual world (seen top right)

Users enjoyed that this flame had an extra layer of physicality to it (the shadow casting), but were disappointed to find that it did not also supply heat in some way as a flame would - they understood why it didn't, but still remarked that "subconsciously" they expected to be able to hold it and feel the effect of the flame on their hand. Although producing heat in this way is very difficult with the technology that is currently available, it does open up the thought process of systems like this to continue down a more multi-sensory approach.

Sometimes physical objects are not directly the best solution

It has become clear within the course of this study that although interactions based upon physical elements offer a great deal of familiarity to users due to their expertise with interacting with the physical world, they also can be a hindrance due to their physical size. For example, if a user has one object that around which they are trying to position Radial Parameter Markers in order to apply some parameter to it, the user can only position a certain number of markers before the workspace is too congested to work efficiently without continuously reorganising these markers. Some of this can be seen as a problem with the Radial Parameter Marker interaction technique, but it is also a reflection upon interacting with physical objects. One of the benefits of a GUI is the ability to modify, and add to, the attributes or parameters of a "virtual" object. These modifications are not embodied in the physical space and therefore we can theoretically add unlimited attributes to an object. With a physical object this is somewhat more difficult: each attribute or parameter may take up additional space. Consider changing the colour of an object in a GUI: a user could right-click a menu option and change the colour, yet with a physical object the user would have to either replace the object with one of a different colour, or have some mechanism to change colour (such as painting the object, or including a colour changing light). Both of these options would require additional space.

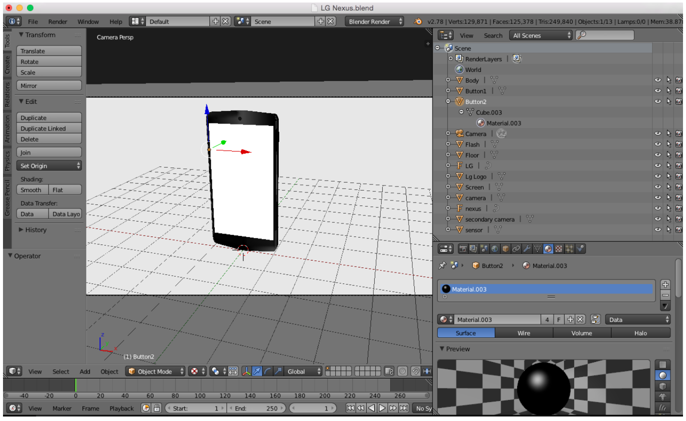

Figure 3.22: A Screenshot of the 3D design tool,'Blender', showing many attributes added to an object. These can be seen to the right hand side.

People act differently when being recorded

During the first round of preliminary user testing, people were invited to interact with the system whilst being filmed by a camera set up to collect video and audio data. However, it quickly became apparent that people would act differently if they were recorded - especially where people had negative comments. In an informal situation these negative comments were openly discussed but when faced with a camera people were less forthcoming. A number of users from the workshop expressed that they felt uncomfortable saying some things if they were being recorded. The process was repeated using a smaller camera [1] (see Figure3.23) with the aim of making data collection less intrusive and therefore make users more comfortable. However, this was not effective since users were still aware that they were being recorded, and thus the same issue arose. The most productive conversations occurred off-camera, and so handwritten notes in diary form were recorded by the researchers throughout the workshop events, which were then written up at the end in addition to taking pictures. This is a technique that is continued through the other research projects for the same reasons.

[1] Small action camera. See https://gopro.com/product-comparison-hd-hero2-hd-hero-cameras

(a) Nikon D3300 **(b)** GoPro HD Hero 2

Figure 3.23: The cameras that was used to record video

Whilst some of the findings discussed here are key to the continued research of this area, others are more tightly coupled with this specific system and are therefore more difficult to generalise. One such finding is again related to the constraints that are implicitly in place on physical objects: they take up space, have weight, e.t.c. In this case, the difficulty comes from the objects not being able to float in mid-air, coupled with the AR marker remote sensing technique that was employed to track objects. In some cases the user would build a hill and wish to place one of the markers on the slope of the hill. When the markers were positioned in this space, the system would find it difficult to detect as it was distorted by the angle of the slope. In this instance the users resorted to balancing the marker on something that would hold it up straight. The marker could now be detected, but it distorted the shape of the hill that they were trying to make. Whilst this is a problem that is tightly coupled to the implementation of the system, it is still useful as a basis for some of the techniques discussed in the following chapters.

3.6 Conclusion

In this chapter I have introduced the first iteration of the design process, as described in Chapter 2; Digital Terra Firma (DTF). The main goal of the DTF system is to synchronise two co-existing interaction spaces: the physical and the digital. Whilst the system does achieve this goal to some extent, the testing that was carried out indicated that a tool such as this relies very heavily on vision as the carrier for information being synchronised from the virtual to the physical (e.g via projection). The concept that the portrayal of this information would be better approached from a multi-sensory perspective is the key outcome of this step, and this idea will be carried forwards into later iterations of the research process. Additionally, large portions of hardware and software that were described in this section are also carried in to later iterations.

Chapter 4

Research Projects: Magnetic Files

4.1 Paper

This chapter was originally written as a paper for academic consumption and presented in San Fransisco, USA at Tangible, Embedded, and embodied Interaction (TEI) (Gullick et al., 2015). It was generally well received, and in addition to being a thoroughly enjoyable and informative experience, the feedback and clarification needed throughout the review and discussion process has positively impacted this chapter.

4.2 Planning

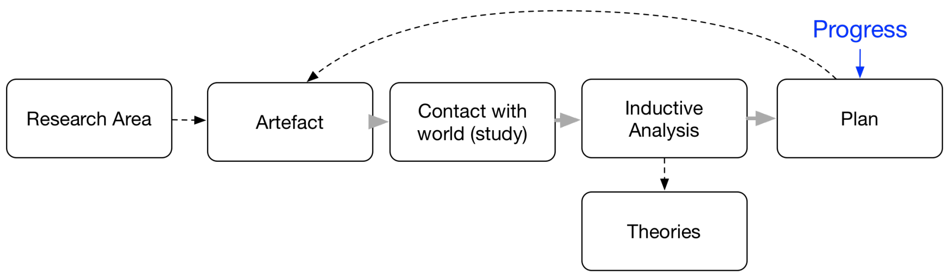

Figure 4.1: Current progress in the design cycle.

The process diagram highlighted in the Methodology Section (Figure 2.1) shows that at this point in the research process, one full iteration though the design cycle has been completed: an artefact has been created though a process of hack & hunch (figure 2.2); the artefact has been connected to the outside world through various user studies; and an inductive analysis has been performed upon the outcomes. Given the knowledge acquired in the form of themes, the research is now at the stage where these findings can be used to help plan the next iteration of the cycle: further investigation of some key areas of intrigue highlighted throughout the DTF project are now possible. The main areas of interest which arose were: the concept of more fully multi-sensory interaction approach to the system; and secondly, the way that physical tools can be used as references to objects, and how we might apply virtual attributes to an object using a physical control system.

To begin looking at this, a small brainstorming workshop was held that involved many of the same participants from the preliminary study performed in the DTF project, alongside some new members. This brainstorming session produced ideas that focused on physical objects that can in some way manipulate many aspects of their physical existence: for example, change their shape, size, colour, temperature, and texture to indicate the status of virtual information. From this session a number of ideas and possibilities for the continuation of this work were produced. The overarching theme of these ideas was to have a number of physical objects that could be 'attached' to a virtual entity and update and change their physical properties to represent the corresponding virtual object.

In a similar fashion to DTF, digital and physical spaces would need to be synchronised to achieve this goal. However, in this iteration of the research cycle more emphasis needed to be placed upon synchronisation across multiple senses, rather than just the visual. As this workshop progressed, it became clear that it was necessary to find scenarios in which lots of virtual information existed that could be expressed physically. Whilst a number of such scenarios exist, it was decided that the most applicable (and easiest to produce) scenario was one of managing computer files. Computer files already contain a huge amount of information - the contents of the file, in addition to meta-data including date and time of last access, last user, and file size amongst many others - and are an easily accessible source of virtual properties that have no inherent physical embodiment. Rather than focus on the traditionally used hierarchical file system that has a more substantial body of academic work behind it, this research focuses on the lesser explored area of “tag based file systems”. These tag based file systems became the basis for much of the interaction explored throughout the rest of this chapter as they may lend themselves to a wider range of interactions that the aforementioned hierarchical model. Tag based file systems (Bloehdorn et al., 2006) are introduced more in-depth in Section 4.2.1.

The previously discussed DTF system was built with the idea of synchronising a real and a virtual space. In reality, it ended up primarily focusing on replicating what is physical within the digital realm, and similarly it focused primarily on vision as the sense with which these two realms were synchronised. MagneticFiles (MF) instead concentrates on the way in which information that is in the virtual but may be hard to represent physically can be fed back to users, expanding to beyond vision and into other senses. In the following sections the MF system is introduced in detail, and the interactions between users and the system are explored.

4.2.1 An Introduction to Tag based file systems

Traditionally, computers store files in a strictly hierarchical format, with each parent folder containing many other child files or folders. This format is one that was initially designed to help in the uptake of personal computers in the early stages of computer systems. This hierarchical format is built to mimic the way in which people would organise paper files into folders in the real world - this mimicry was part of the “Desktop Metaphor”: a metaphor that allows users to more easily understand and predict the actions of their computer system. As a result, a modern digital hierarchical file system would be relatively simple (yet physically enormous) to replicate in the real world: each file could simply be replaced with a real paper document, and each folder with a real world folder. However, hierarchical style file systems are not the only file system formats to exists for a number of reasons: in addition to a number of low level computing issues, they also force users to remember the exact location of their files, and make difficult decisions about where to place their files (Karypidis and Lalis, 2003). Additionally, hierarchical structures are not flexible enough to accommodate content that overlaps two folders that are stored in discrete parts of the hierarchy.

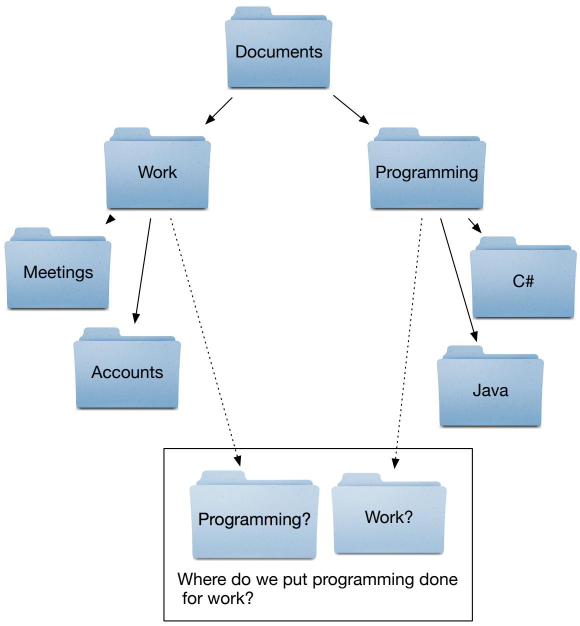

Figure 4.2: Hierarchical file systems can limit our data storage decisions.

As circumstances change, a hierarchy of files may not reflect the content in a way that is helpful to the user: consider the scenario highlighted in Figure 4.2 where there are multiple suitable storage locations for one file. As a result, many people choose alternatives to this hierarchy, such as using a search tool, to manage their files. This has lead to alternative file system paradigms being explored by the research community and leading software companies. One of these alternatives is the “tag-based” file system (Agarawala and Balakrishnan, 2006, OSXTags). Put simply, tag based file systems allow the user to keep their files in a large pool, attaching keyword tags to make each file unique.The user can then search this pool by specifying a number of search tags, eliminating irrelevant files from the search. This alternative view on file systems seems particularly suited to our needs as it focuses on tags or attributes of files for manipulation, and is largely unexplored with relation to TUIs. Tag based file systems are therefore the underlying virtual element upon which the interaction is built within this system.

4.2.2 The Magnetic Files concept

If we combine the ideas behind tag based file systems and the previous concepts discussed of physical/digital hybrid space (consider the interactions of “MusicBottles” discussed in Chapter

1), a physical pool of files can be imagined as existing within a physical, three-dimensional space. These files could be manipulated using physical means: they could be separated into categories by hand; or collected in a container and moved between different places. Many of their physical properties would be experienced by a user interacting with them: if a file is picked up its weight, temperature, texture, smell, and rigidity can be evaluated along with any number of other physical properties. For example, in this scenario it could be possible to physically evaluate the size of a file on disk by feeling its weight, or observing its physical size.

The idea of the Magnetic Files project is to come somewhat towards achieving this scenario - or as close as possible given the limitations of technology, time, and money available for this work. Rather than aim to achieve this in a three-dimensional space, it seems more applicable to mimic this in a two-dimensional space: initial research into available technologies highlighted numerous problems with attempting this three-dimensionally, largely due to limitations on technological availability and cost, but this problem is less apparent with two dimensions. In two dimensions the concepts investigated will be the same, but the technological complexity is decreased substantially. Utilising much of the same hardware as described in DTF it is possible to project something representing files on to the workspace. However, the system must still have some way of injecting physical properties into the space in which the projected file exists, and users must have some way of interacting with them.

4.3 Artefact Creation

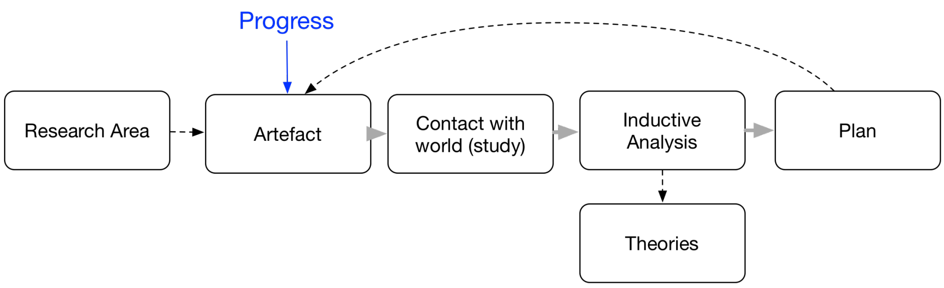

Figure 4.3: Current progress in the design cycle.

In order to achieve this goal two possible options for the physical build were investigated. Firstly, a table (as the 2D plane) that has various embedded electronic components was considered. These components would allow the table to change its properties in any area. For example, it could create localised heat/roughness on the table. The physical properties could then be programmed to follow around the files as they are projected on top of the table.

Figure 4.4 shows an early Computer-Aided Design (CAD) sketch outlining one possible way that the table could have been constructed. It shows a table with linear actuators controlling rods in much the same way that shape changing tables have been constructed in the past, such as with the inFORM project Follmer et al. (2013). In inFORM, actuated rods are moved up and down on a table's surface to mimic a shape changing device. Whilst the goal with MF is not to change the shape of this surface, a similar technique can be used to introduce rough surfaces to the table. The elements (rods,Light Emitting Diode (LED)s, nichrome wire) shown in the sketch are controlled by micro controllers and can raise or lower a rough surface (such as a course grit sand paper) to change the roughness of the table at the relevant spot. Nichrome wire is also run over the surface allowing the surface to also be heated, controllable by area. Similarly LEDs are inlaid into the surface of the table (flush with the surface) so that extra lighting information may be passed onto the table.

Unfortunately, this concept was dismissed as not ideal for this work due to large hardware costs incurred from so many required components. Similarly, the complexity of the hardware required to drive and coordinate these components would require complex software to control. Given a table of two meters by one meter as shown in the sketches above, approximately 200 linear actuators, strips of nichrome wire, and RGB LEDs would be needed if the patches they control were 10cm x 10cm in area. An Arduino or equivalent micro controller can control 3-4 linear actuators at a time so more extensive control systems would be needed to coordinate multiple micro controllers. This increase in complexity would also result in a more awkward space within which to interact. For example, any objects placed onto the surface block the surface of the table, meaning that, for example, sand could not be used on top of the surface such as with DTF.

The second option, which was chosen for this study, was to construct some form of robot that could be remote controlled to move around the table. Much like the table described above, it modifies its physical properties, but the number of components needed is greatly reduced as the robot could move to where was required. The cost of covering a table in the required electronics would have been at least an order of magnitude higher than that of constructing a small number of robotic actors. A table could arguably be more seamless in terms of a user

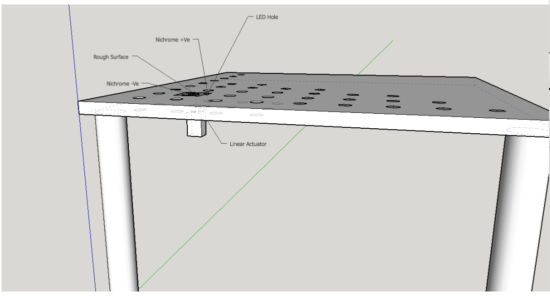

(a) Side

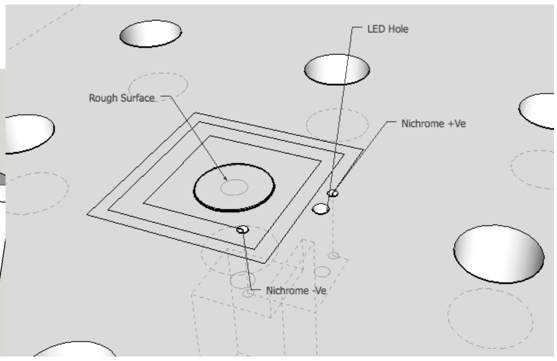

(b) Top

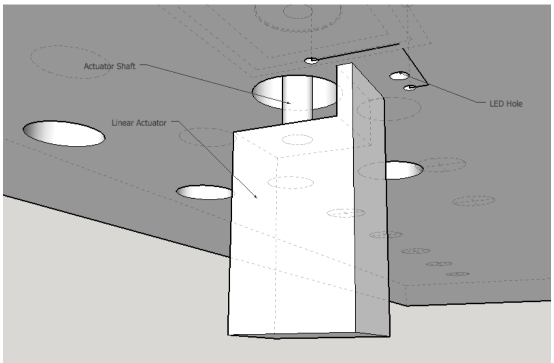

(c) Bottom

Figure 4.4: An early idea for an attribute changing table

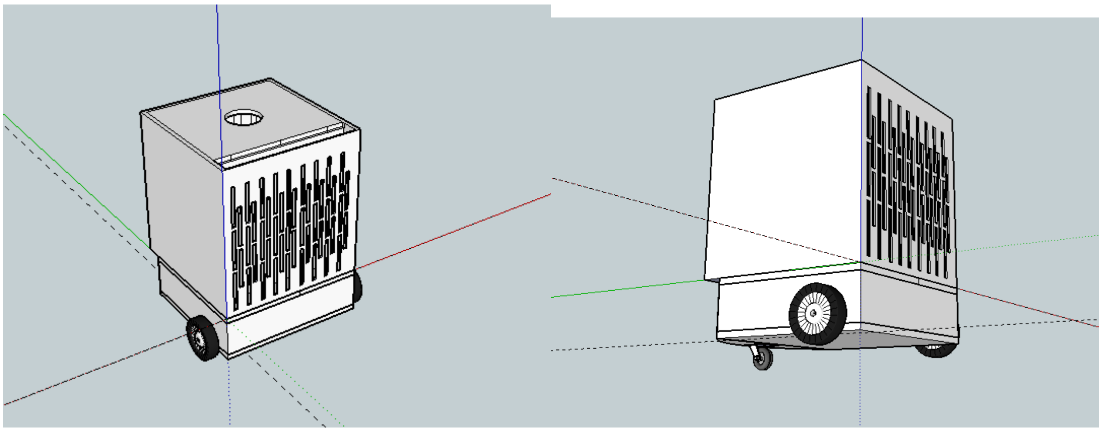

(a) Filebot sketch view 1 (b) Filebot sketch view 2

Figure 4.5: A Sketch of an early idea for the Filebot

experience, but it does not allow certain attributes, such as weight, to be expressed: a user cannot pick it up and manipulate it.

As a result, the decision to build a small number of robots to represent files in a tag based file-system was reached. These robots were later called "Filebots", and they were designed to operate on a flat tabletop surface. Large portions of the DTF system were re-used throughout this iteration of the project: the projector, depth sensor, and camera were set up as they were previously and allowed depth and colour mapping of the space. This set up allowed the computer to build up a virtual understanding of what is happening physically in the interaction space, and augment it digitally using the projector. The use of AR markers as recognisable tags within the space were once again employed to allow easily manipulation of various elements of the system.

4.3.1 The Filebot Chassis

These Filebot construction was split into two phases. The first of these phases was focused on what was termed the "chassis" (the lower part in images shown in Figure 4.5. The upper part is discussed later). The Filebot chassis has two main functions: to move the Filebot about the surface (and as such it must house motors, batteries, controllers and various other electronic components); and to facilitate detection and control by the computer running the MF software.

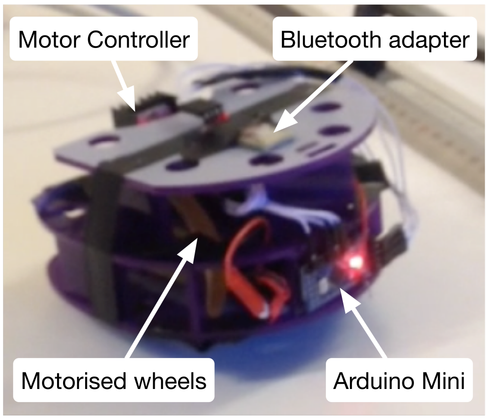

Figure 4.6: Version 1 of the Filebot Chassis

Two major iterations of the Filebot chassis were trialled. The first was a simple laser cut chassis which used interlocking tab and slots to fit together[1]. This chassis housed all of the electronic components to control two large motors and gearboxes. The way these electronics modules fit together is shown in Figure 4.7.

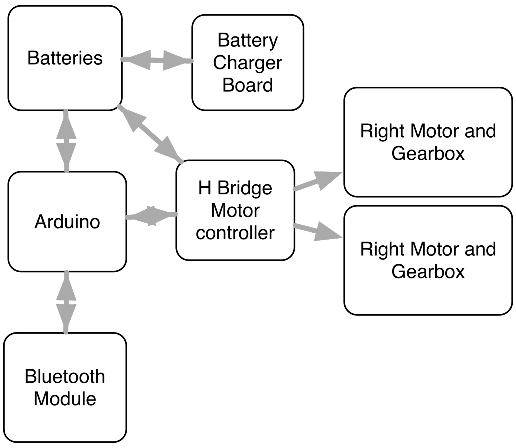

Figure 4.7: A diagram illustrating how the Filebot Chassis electronics fit together.

The first iteration of the chassis housed a number of electronic components which were laid out according to Figure 4.7. Figure 4.8 shows the Arduino pro mini microcontroller board used to control each of the components. This microcontroller could be programmed with a micro USB cable and the Arduino Integrated Development Environment (IDE). In essence, this

[1]Based on the Hack-e-bot available here: https://www.thingiverse.com/thing:166465

microcontroller board communicated with the computer and controlled the local components as dictated by the computer. Everything on the Filebot was powered by two 3.7 volt Lithium Polymer (LiPo) batteries which are also shown in Figure 4.7.

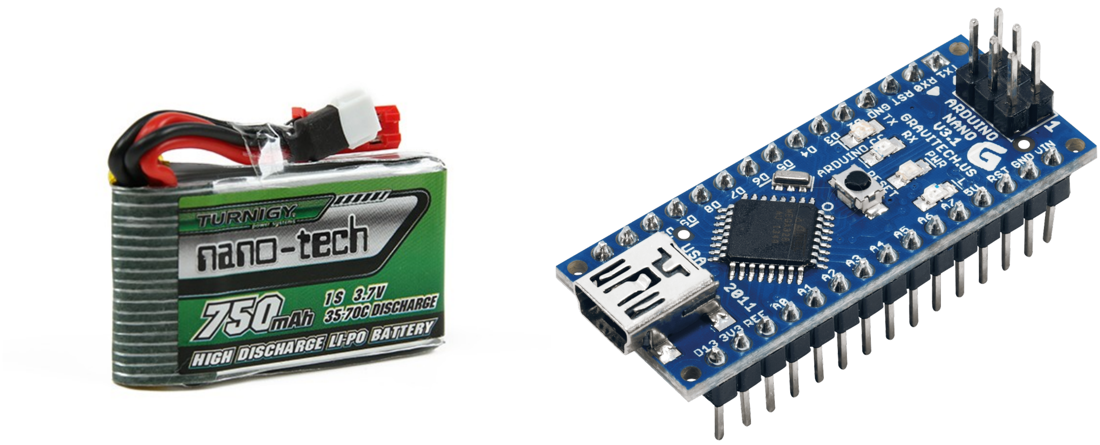

(a) 3.7v Lithium polymer Battery

(b) Arduino nano

Figure 4.8: The main electronic components to the V1 Chassis part 1

This LiPo battery supplies power at 3.7 volts, whereas the components used inside the Filebot needed high voltages. To allow this, Buck Boost voltage regulators such as the one shown in Figure 4.9a allowed this voltage to be boosted as required. LiPo batteries require specific charging profiles, so a module to charge the batteries via USB was also included, as shown in Figure 4.9b. The batteries allowed each Filebot to be powered for approximately two hours between recharges, depending on what the Filebot was doing.

(a) Voltage Regulator (Buck/Boost)

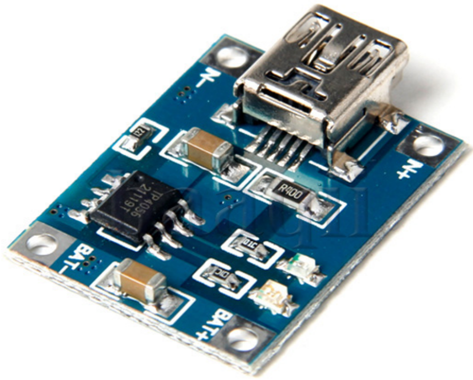

(b) Battery Controller Board

Figure 4.9: The main electronic components to the V1 Chassis part 2

To provide movement, the Filebot Chassis housed two separate motors, one for the left wheel, and one for the right. Steering was achieved using 'skid-steering': a technique that modulates the rotational speed of each wheel to achieve the required turn. Each of these motors was connected to the wheel via the gearbox which allowed the wheels to turn more slowly, but with increased torque. As the microcontroller is only able to output between 0-5 volts, turning the motor backwards was not possible without the use of a H-bridge module. The H-bridge takes two signals: the first is 0-5 volts to indicate speed to turn; the second is 0 or 5 volts to indicate forwards or backwards. The motor, gearbox, and H-bridge board are shown in Figure 4.10.

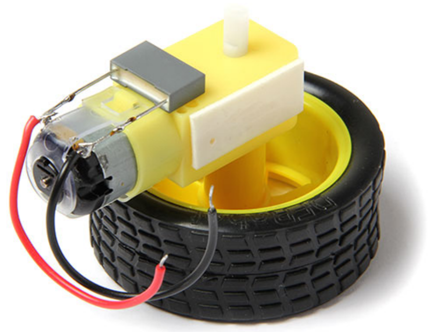

(a) Motor & Gearbox

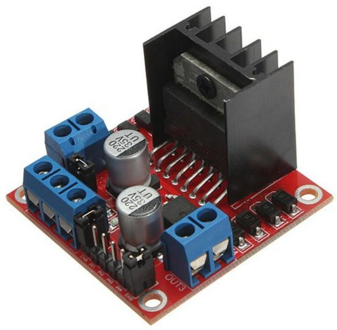

(b) Motor Driver (H Bridge)

Figure 4.10: The main electronic components to the V1 Chassis part 3

The communication between each Filebot and the computer was handled via the Bluetooth module shown in Figure 4.11. Once connected to the Arduino, it allows communication between any Bluetooth device and the Arduino as if plugged in via USB. This communication was in the form of a custom command protocol which is discussed later in this chapter.

(a) Linvor Bluetooth Module

Figure 4.11: The main electronic components to the V1 Chassis part 4

Whilst this version of the chassis was fit for purpose in the fact that it could move and be controlled by the computer, it had a number of small issues that needed to be fixed before the Filebots worked as intended. Firstly, due to the thickness of the cut produced by the laser cutter, the tab and slot joints did not hold completely rigid and glue was needed to rectify this. However, this process of gluing meant that it was difficult to adjust the electronics inside the chassis. The gearboxes also had some slack between the internal gears which meant that turning was not as precise as needed for accurate control using the computer. Additionally, issues that stemmed from wheel slipping during skid-steering contributed to errors tracking the Filebots. To address these issues, a second iteration of the chassis was constructed, based heavily upon the Zumo robot chassis[1].

This second iteration improved upon the first by being much smaller, and including a gyroscopic sensor alongside a compass and rotary encoders. The gyroscope and compass combination allowed the robot to detect its orientation in free space. Rotary encoders accurately tracked the rotation of each wheel to allow the precise control of the wheels for accurate turning. These additions meant that alongside being a much more compact robot, the control offered by the second version of the chassis was much greater than the first iteration. Figure 4.12 shows the second version of the chassis.

[1] https://www.pololu.com/product/2506

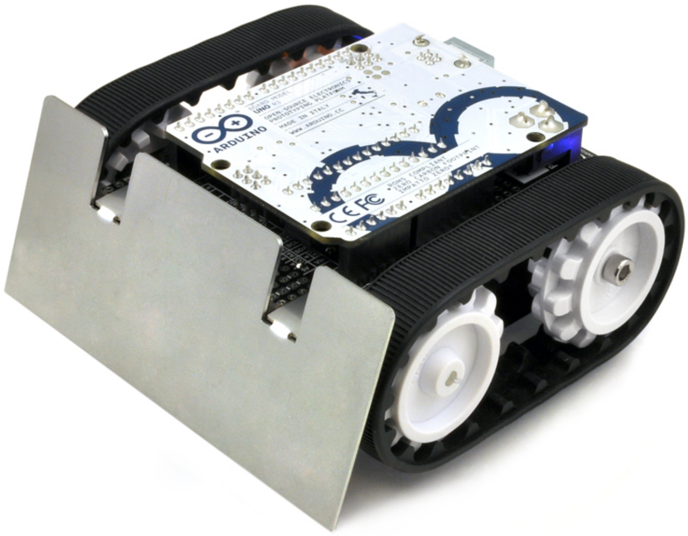

Figure 4.12: The Zumo robot that performed the task of the second iteration of Filebot Chassis.

Many of the components were replaced from Chassis version 1 as they were already incorporated into the Zumo controller board. The Bluetooth module was still used to allow communication between the robot and the computer. The LiPo batteries and charging board were replaced with four AA batteries. However, The Buck Boost control board was still used to power some of the other modules discussed later in this chapter.

4.3.2 Detecting and control of the Filebots

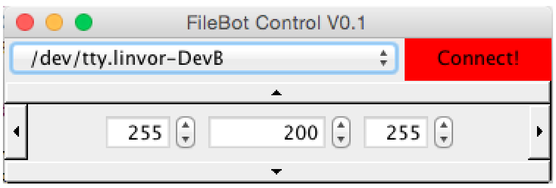

Figure 4.13: An early version of the remote control software for the Filebot

The initial version of the chassis was remote controlled over Bluetooth via a small Java application that moved the robot respective to arrow key presses on the computer keyboard (a screen shot is shown in Figure 4.13). Later versions used automatic detection and path finding for the Filebots. Although tracking of Filebots was initially attempted using only the positions reported from the rotary encoders on the Filebots, this proved inaccurate (presumably due to slippage in skid-steering). As a large part of the development of the DTF project focused around the inclusion of AR markers, they had already been found to be easily detectable by the camera incorporated into the Microsoft Kinect. Accurate location tracking was achieved by attaching an individual AR marker to act as an identifier. Through the use of the nyARToolkit library, access to position (x,y coordinates), and also to orientation/rotation for each AR marker is available. Given this data, the computer worked out the best path between where the Filebot was, and where it needed to be (rotation and distance), and relayed the commands to move the robot. The protocol is discussed in more detail later in this chapter. A still frame from this process can be seen in Figure 4.14.

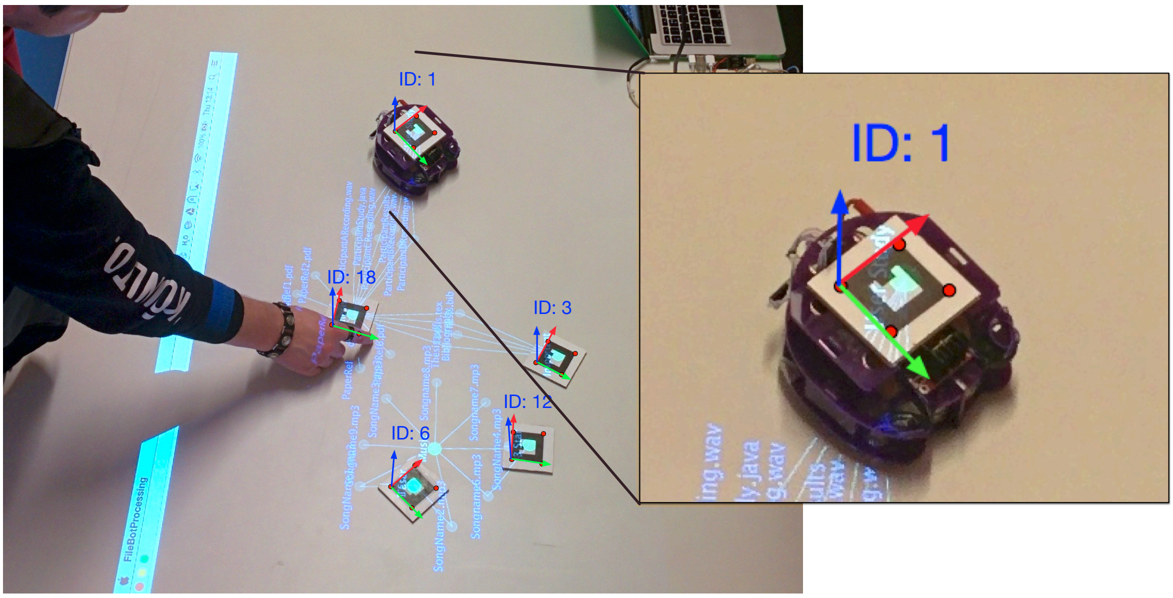

Figure 4.14: A screen shot from the AR Detection viewer software during testing with Chassis V1, showing AR Detection debugging markers.

Each AR marker was detected using its four corner points. Given that it is known to be square, three-dimensional positioning and orientation could be calculated. Similarly, each marker is unique, so the ID (and corresponding object) could be determined. The algorithm to calculate the movement path needed by the Filebots to move from their current location to their required location was simple: they first rotated by the necessary number of degrees to be able to face the new target location, and they then drove forward the required distance, at all times receiving new instructions from the computer over Bluetooth connection. For the purposes of this work this algorithm works fine, but more complicated path finding algorithms would be needed if many Filebots were in a workspace so as to avoid collisions and overlapping paths.

4.3.3 The Filebot 'BackPack'

A box that housed all of the electronics and components that allow the Filebot to change its physical properties was attached to the top of the chassis, which was later called a 'Backpack'. If the digital attributes that a file has (and that the Filebot might need to represent) are considered, we can produce something similar to Table 4.1:

Table 4.1: Physical properties mapped to digital file properties

Property	Information
Shape	File Type
Size	Perceived Importance.
Weight	Word Count or File size.
Texture	Draft Status of file - Rough vs Final.
Rigidity	File type or word count.
Temperature	Time since last modified.

Whilst there are many potential ways that these different attributes could be controlled electronically, the following were trialled for the purposes of the Filebot.

Temperature

In order to change temperature, Nichrome wire was used. Nichrome wire is a metal wire which is used much like filament in a light bulb: passing a current through Nichrome wire will increase the temperature of the wire. This property was used in early versions of the backpack by embedding the wire in the outer material, which could then be heated as needed. Alternatively, a Peltier element could be used: it is physically less flexible that Nichrome wire, but allows both heating, and cooling depending on the direction of the current passed through it. After trialling both approaches, a Peltier Element was used. This can be seen attacked to part of the backpack in Figure 4.15.

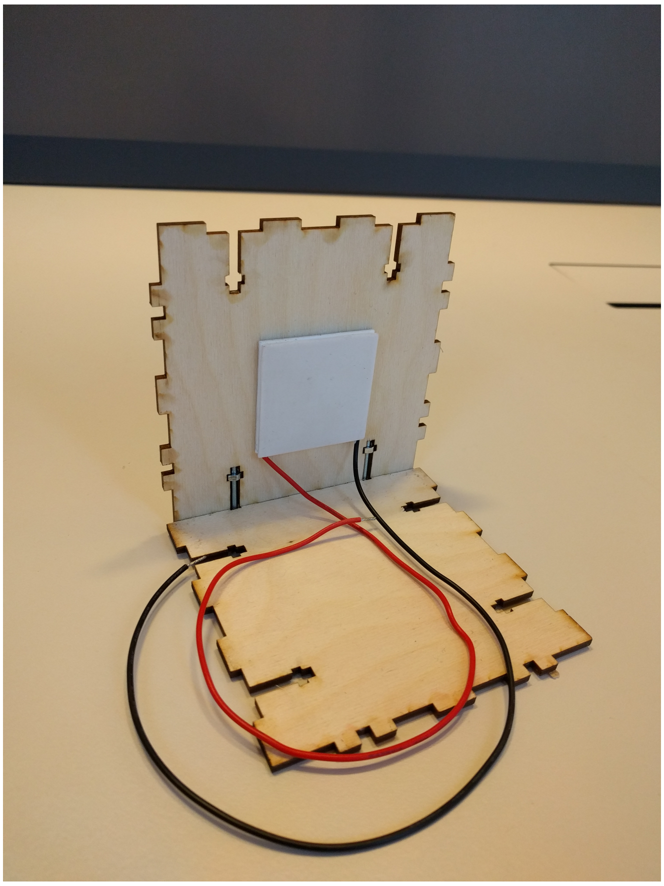

Figure 4.15: An image of the Peltier Element as used in the Backpack.

Texture

A number of different methods were trialled to adjust the texture of the Filebot - from rotating a panel with different grades (roughness) of sandpaper to raising many small spikes using servos. A servo (short for servomechanism) is a mechanical device that can be electronically instructed to achieve a certain position which is typically achieved with the use of motors and some positional feedback mechanism. Eventually, the idea to twist tightly stretched Lycra using a servo produced more reliable results that raising spikes. When stretched tight and untwisted (like the top of a drum), Lycra presents as a smooth flat surface. However, when it is twisted at the centre by the servo, the Lycra wrinkles creating a rough surface of valleys and troughs in the material. This part of the Backpack can be seen attacked to part of the backpack in Figure 4.16.

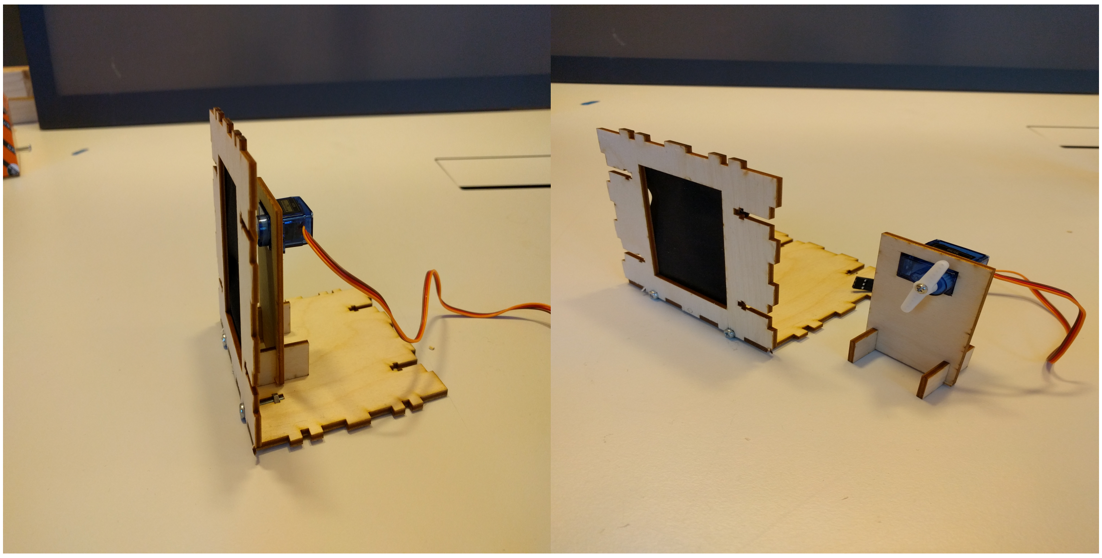

(a) Separated **(b)** In Position

Figure 4.16: An image of the Lycra twisting mechanism as used in the Backpack.

Rigidity & Weight change

Rigidity or weight of the robot could theoretically be changed with the inflation of an internal balloon using an air pump, or internal reservoir being filled with compressed gas or a heavy substance like mercury (see (Niiyama et al., 2014)). Weight change was difficult to achieve satisfactorily as it required the use of external reservoirs of material to be transferred to the Filebots. This would require pumps and flexible tubing between the reservoir and the Filebots which was not ideal. Weight change as a changeable parameter was therefore discontinued. Despite this, rigidity control was still achievable without the user of an inflatable bladder in a similar fashion to texture control using Lycra. In this case, rather than twist the Lycra using a servo, the Lycra is stretched by a rotating lobe on a shaft, which is once again activated with the use of servos. This part of the Backpack can be seen attacked to part of the backpack in Figure 4.17.

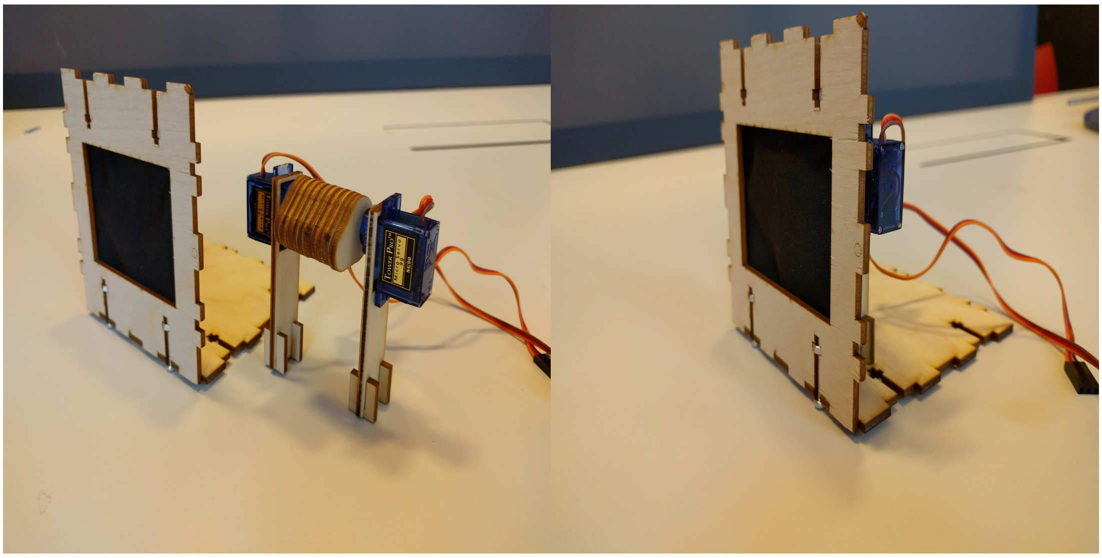

(a) Separated **(b)** In Position

Figure 4.17: An image of the Lycra stretching mechanism as used in the Backpack.

Colour

The colour of the Filebot was controlled via RGB LEDs. These LEDs can produce any colour, which is controlled via signals from the microcontroller. These LEDs were used to illuminate a frosted panel on the side of the Backpack. This part of the Backpack can be seen attacked to part of the backpack in Figure 4.18.

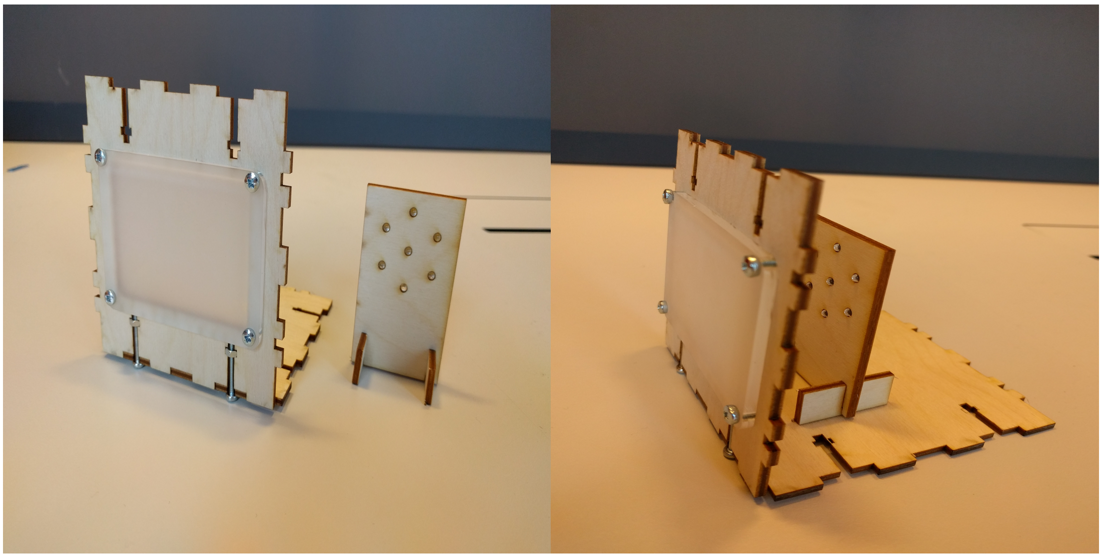

(a) Separated **(b)** In Position

Figure 4.18: An image of the LED mechanism on the Backpack.

Once assembled together, these different modules construct the complete Backpack which could be connected to the chassis. Once attached, the Backpack connected to the microcontroller in the chassis and allowed the computer to control the physical attributes of the Filebot via the Bluetooth link. A custom communication protocol was used over the serial link provided by the Bluetooth to control various attributes (see Section 4.2 for a more detailed look into this protocol).

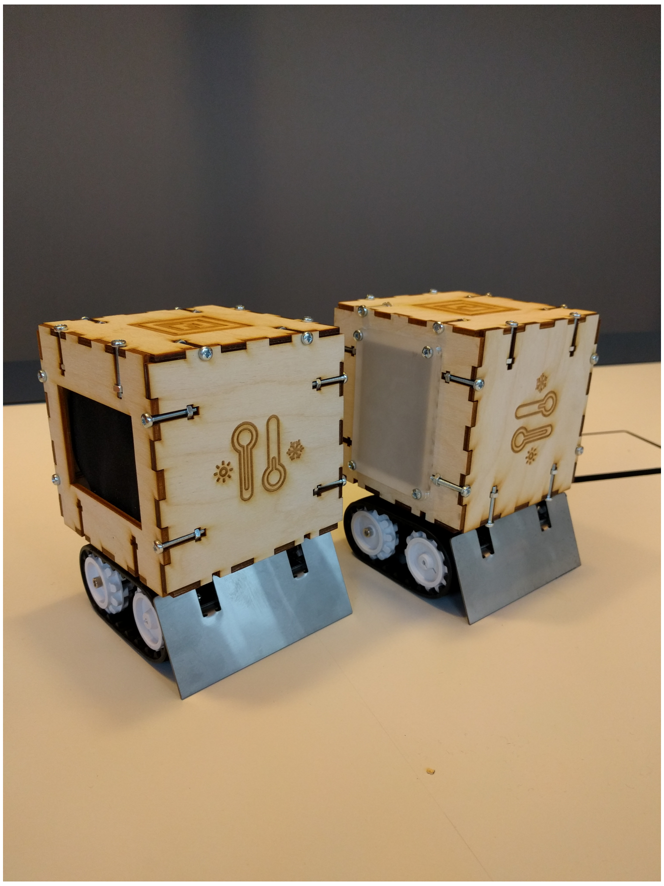

Figure 4.19: The assembled Backpacks on top of Chassis v2

While multiple working version of the complete chassis were trialled throughout this build, only two fully working versions of the backpack were constructed for the Filebots due to time constraints. However, multiple versions of the mechanisms used in the backpacks were trialled. The final version of the backpack is shown in Figure 4.19. A closer look at the Backpack from different angles is shown in Figure 4.20. It comprised a laser cut wooden cube, again using tab and slot fittings, that attached to the top of the Filebots to allow physical property changes. The top face of the cube was taken up with the AR marker. The remaining sides were taken up with a temperature controlled face (controlled with an internal Peltier Element), a texture controlled face (Stretched Vinyl), a light controlled Face (RGB LEDs), and finally a face that can bow in or outwards to change rigidity (Lycra stretch controlled via internal servo).

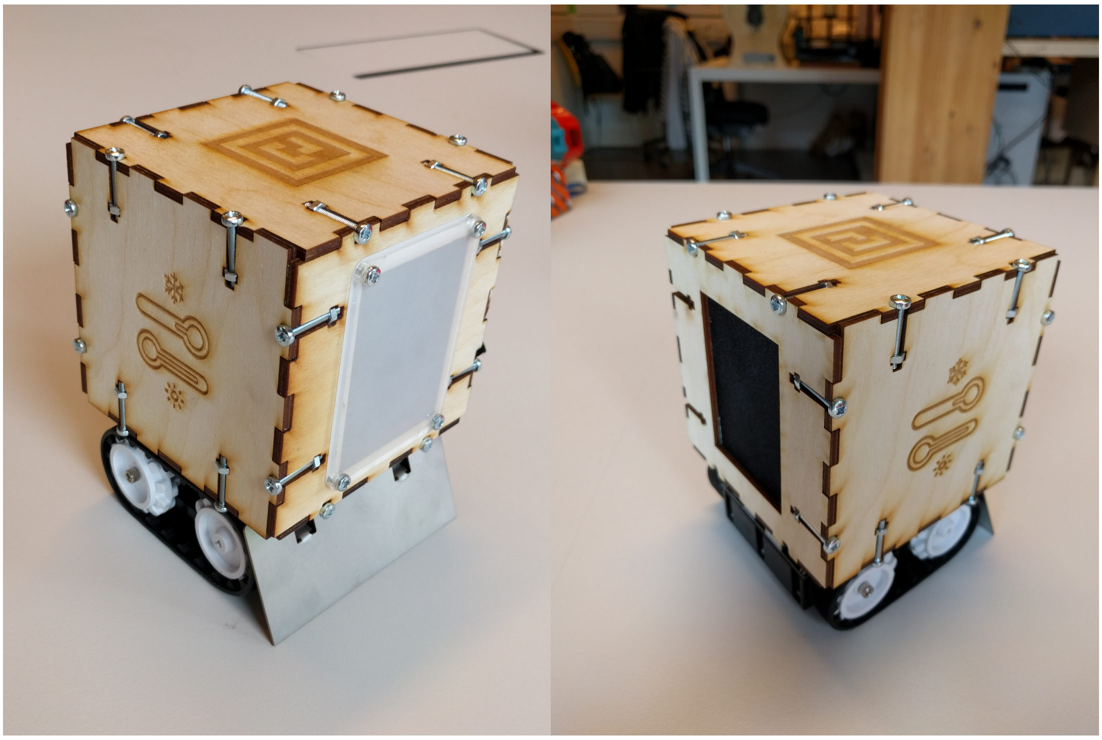

(a) BackPack close up 1.

(b) backPack close up 2.

Figure 4.20: Close up view of the backpack

All together, the Filebot could be tracked, turn and move, and update a number of its physical properties as directed by the computer. This was all controlled by the software running on the computer with the camera attached nearby, and is discussed in the next section.

4.3.4 Software

The software the drives the Magnetic files is a layer built on top of the original software used in the DTF system. A number of extra software components were added.

The first of these software components is a module that interacts with the user's file system. The computer that this system runs on uses the Hierarchical File System Plus (HFS+) file system (the default file system for Mac OSX computers at the time of writing, a replacement and improvement upon the previously used Hierarchical File System (HFS)) - which supports tag based files. When directed to a folder, this module gives the program access to all of the files and their respective tags and attributes contained within the designated HFS+ directory. For testing a directory was set up with approximately 1000 files in it. These files were designated as a mix of files from a work like environment (spreadsheets and documents) and a home like environment (video, music, photos).

A module which handled the representation of these files when projected on to the work surface was also added. It projected every file as an element on to the work surface. The locations of these files were determined by a very simple pseudo-computerised physics engine (based upon JBox2D[1]) that initially mimicked the way in which physical objects would float around in zero gravity, but was later used to apply forces to separate the files, such as repulsion, spring, and dampener connections.

A module that handled communication with the robots via the previously discussed Bluetooth module was employed. This module took input from the Kinect's RGB camera to determine each Filebot's location and direction, and calculated the rotation and path between its current position and the target position (the location of the projected file with which it was coupled). The appropriate commands were then sent to the Filebot via a simple protocol over Bluetooth serial just for this purpose, which is described in Figure 4.2:

[1] http://www.jbox2d.org/

Table 4.2: The Filebot Serial control protocol.

Controls	Command	Comment
Set motor speeds	speed [left,right] set [value];	Moves the Filebot.
Set temperature	temp set [value]:	Sets the temperature.
Set LED	filebot led set [r][g][b];	Sets the LED colour.
Set Texture	filebot texture set [value[;	Sets the texture.
Stop	filebot stop;	Stops the Filebot.
Reset	filebot reset;	Resets the Filebot.
Get	filebot [speed,temp,led,texture] get;	Queries a value.

This module also handled the synchronisation between the properties of the file and the electronics on the Filebot. Upon execution of the program, this module needed some configuration: it required a look-up table of how each physical property of the Filebot should be mapped to each file property (e.g. file size was represented by temperature of the Filebot). This is supplied by editing an eXtensible Markup Language (XML) configuration file that accompanied the program.

Lastly, a module that handled the interactions of the system was used. This module handled the detection of AR markers, user input, and updating the other modules accordingly. These interactions are described further in the next section.

Whilst the technology used to create the system has been covered, the interactions with this system have not been discussed in detail. These interactions can be considered as existing within two strands of parallel work. The first of these strands investigates how a person can interact with so many files in a small space: this focuses on the navigation around the interaction scenario specific to MF (tag based filesystems), and so potentially has less relevance to the overall goals of this work. The second of these focuses on how the Filebots factor in to this interaction, and what methods would need to be employed to stretch a relatively small fleet of Filebots to represent a larger number of files.

4.3.5 Interaction 1: The Magnetic aspect & Magnetic Tags

Throughout a number of preliminary testing, work-shopping and brainstorming sessions, the way that a pool of tagged files could be searched was explored. One of the most successful concepts that tied in with the goal to mimic a physical system was to represent the files as existing within a physics engine. A physics engine is a software component that has embedded into it the rules that govern the physical world. Any virtual objects that are governed by this physics engine are subject to similar physical phenomena as the real world: they can collide with one another and are effected by gravity, aerodynamics, and other similar rules that are apparent in the physical world. Using a physics engine, the files projected can be subject to these same rules in such a way as they mimic real life and as such, come some way towards meeting the goal outlined at the start of this chapter. Treating these files as as "magnetic", and manipulation via "magnets" was introduced as a mechanism of sorting the files as it allows interaction with the projected files in a way that mirrors a physical process. The concept works by treating the attributes of a file, and objects that we call "Magnetic Tags" as magnetically attracted to each other if they match. Magnetic Tags are physical objects that can be manipulated on the workspace, and are simply identified by a marker (AR marker, as discussed with reference to DTF). Each "Magnet" has a single associated tag, and the user has access to many Magnetic Tags at a time. When placed, the system detects a Magnetic Tag, and each file visualised in the pool responds accordingly. Files with a matching tag are attracted towards the Magnetic Tag, and those without a matching tag are repelled. In some cases, this attraction is not only a binary attract/repel property, instead strength of the attraction is dependent on the degree to which the Magnetic Tag matches the file tag. For example, the "Good Music" tag will attract all rated music files, but will most strongly attract "5 star" music, and least strongly attract "1 star" music. For a clearer description of the Magnetic Tag concept please see Figure 4.21 which shows a conceptual diagram, and 4.22 which shows this same scenario performed on the system.

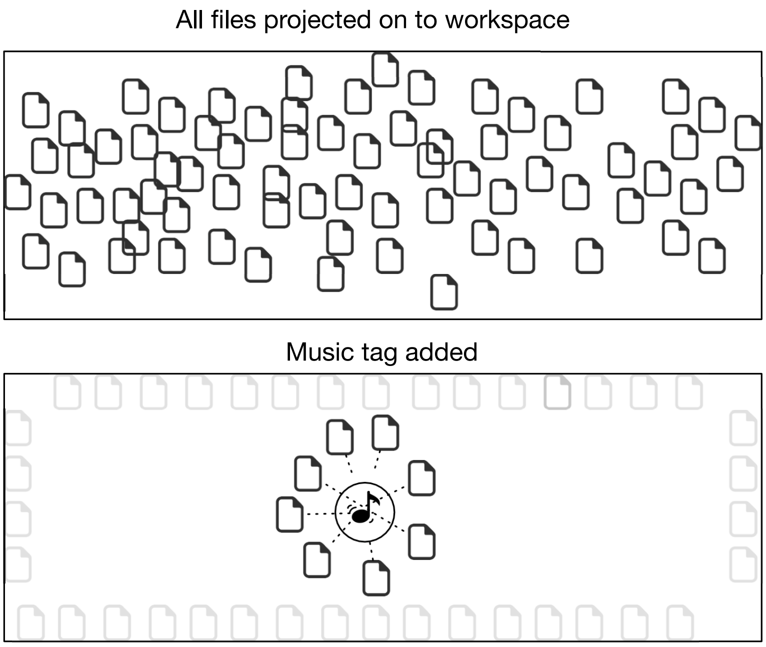

Figure 4.21: A diagram showing the MF concept when a 'music' tag is introduced to the pool of files.

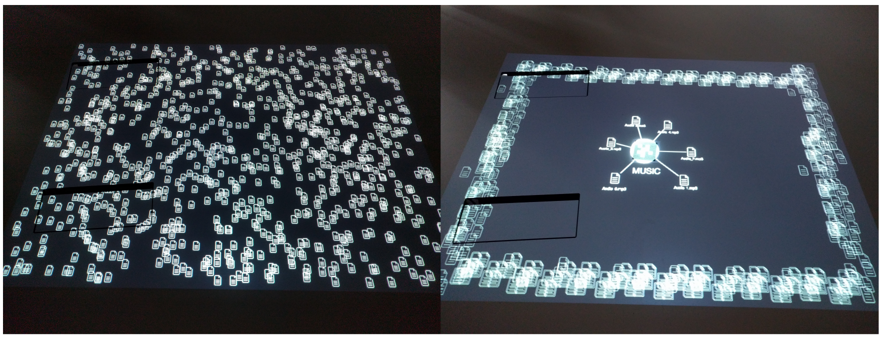

Figure 4.22: A diagram showing the actual MF system when a 'music' tag is introduced to the pool of files.

During typical usage, a user will place multiple "Magnetic Tags" onto the workspace simultaneously, resulting in localised grouping of files with similar tags. Using "Magnetic Tags" as a file search tool allows the user to search through their pool of files, visualising it much like a Venn diagram [1]. This illustration shows the "Coding", "Thesis" and "Academic" Magnetic Tags placed on a surface, and the corresponding files arranging around these tags: their positioning

[1] a diagram representing mathematical or logical sets pictorially as circles.

appropriate to the Magnetic Tags to which they are attracted. These can be seen in concept in Figure 4.23 and reality in Figure 4.24.

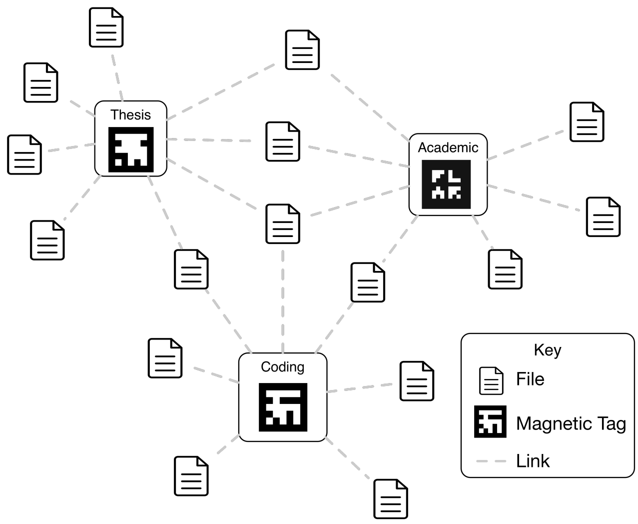

Figure 4.23: An illustration of the system when multiple magnetic tags are introduced to a pool of files

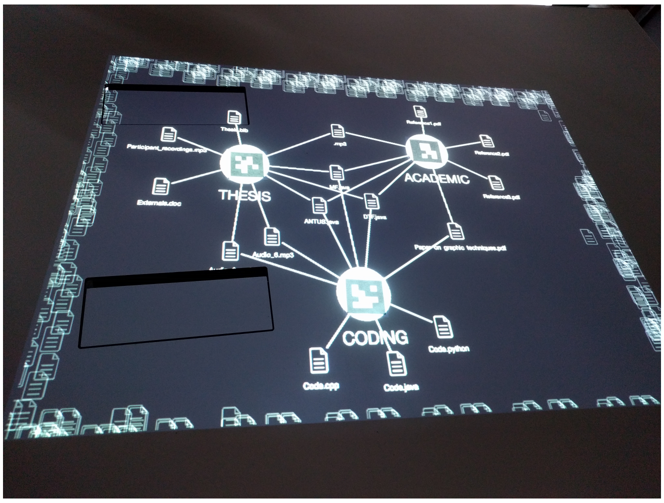

Figure 4.24: An image of the system when multiple magnetic tags are introduced to a pool of files

Using this technique it is possible to tangibly explore a vast number of files and quickly and efficiently. Whilst the workspace initially looks messy due to the vast majority of files displayed (see first image in Figure 4.21), by reducing the opacity of files that are irrelevant to the current search (i.e not attracted to any Magnetic Tag), it is possible to limit the amount of information that is being projected on to the workspace and presented to those using the system as shown in the figures above.

At this point however, these files are only projected onto the table. They offer little in the way of physical interaction other than the apparent attraction and repulsion to the Magnetic Tags. Section 4.3.6 introduces the inclusion of the Filebots previously discussed into this interaction to provide these extra physical feedback systems.

4.3.6 Interaction 2: Expanding the interaction with FileBots

So far, the physical build of the Filebots has been discussed, but not how they fit in with the overall MF system interaction which is the goal of this subsection. In an ideal world, it would be feasible to produce an infinite number of Filebots, and they would be tiny in size: this way, no matter how many files the user was exploring using Magnetic Tags, there would be enough Filebots to represent every file in the space. Unfortunately this was not the case as only two Filebots were created. In the following few paragraphs the 'ideal' scenario will be used when discussing how the Filebots are intended to work, and afterwards there will be a discussion of how something close to this was achieved with so few Filebots.

The concept for the combination of Magnetic Tags and Filebots is that as the user navigates around the pool of files using Magnetic Tags, a Filebot "attaches" itself to a file: one Filebot per file. As the projected files are moved around the space by the users interactions, each Filebot mimics the movement of its paired file so that it is always positioned identically to the projected file. If the projected file is no longer relevant to the users search, the Filebot removes itself from the search area (and the projected file stops being projected). When a new projected file becomes relevant to a search and is projected onto the system, a Filebot once again attaches itself to the file. As the user navigates through their pool of files, they are causing Filebots to move around the space - essentially acting as a physical proxy for that file. The user can see many of the properties of that file and the Filebot adjusts its properties as discussed previously to best represent the file to which it is attached. Those properties that cannot be seen (e.g temperature) can be perceived by the user in other ways, such as through touch. Figure 4.25 shows a sketch developed in the Hack & Hunch stage of this project that depicts this scenario. Using similar tags to those in 4.23, the image shows each file coupled to a Filebot, which modifies its physical properties to best represent the file. Figure 4.26 shows a similar scenario of the real system.

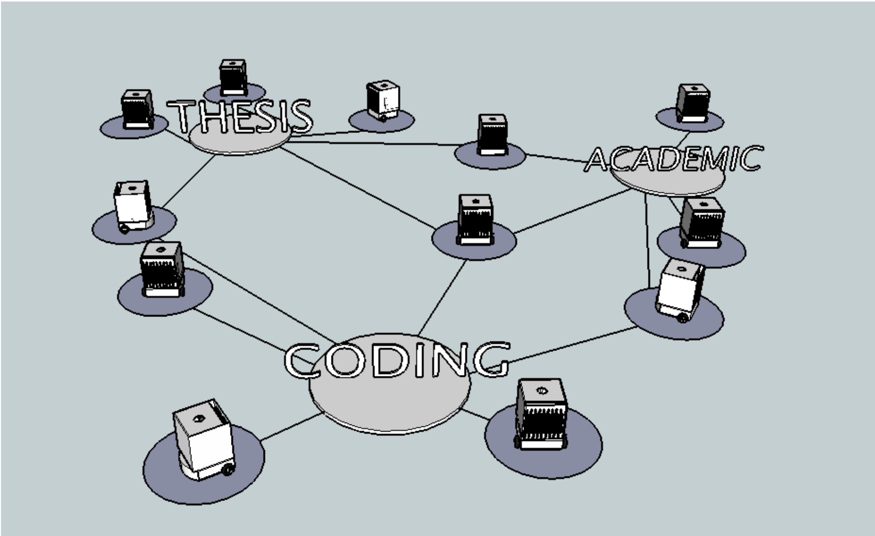

Figure 4.25: An illustration of the tool when multiple magnetic tags are introduced to a system in conjunction to Filebots

Due to technology and time limitations it was not feasible to create more than a small number of these Filebots, and the Filebots created were larger than ideal. Two Filebots were created: one to prototype each chassis design. To make the system work with such a small number of robots, an area of the space was designated as the 'bench' (similar to the bench used in DTF for unrecognised AR markers) in which either robot could wait whilst it was not in use. When a new robot was needed it would leave the bench and 'attach' itself to a file. As their were often too many files for the two Filebots to represent, a researcher was always present at demonstrations who could re-task the Filebots with a different virtual file as the users requested. Whilst it can be argued that this would negatively effect the quality of the discussion and interaction people had with the system, these Filebots and Magnetic Tags were used as props to provide context and inspiration for discussion and as such the discussions were still rich with ideas and insight. This is discussed further in Section 4.4.

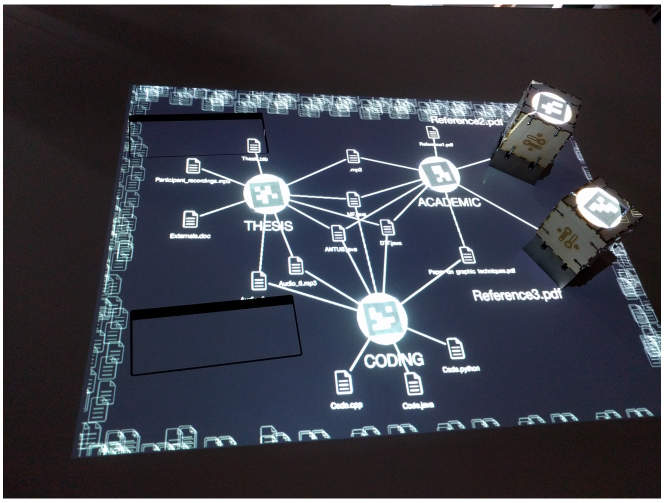

Figure 4.26: A picture taken when Filebots are used in conjunction to Magnetic Tags.

4.4 Contact with the World

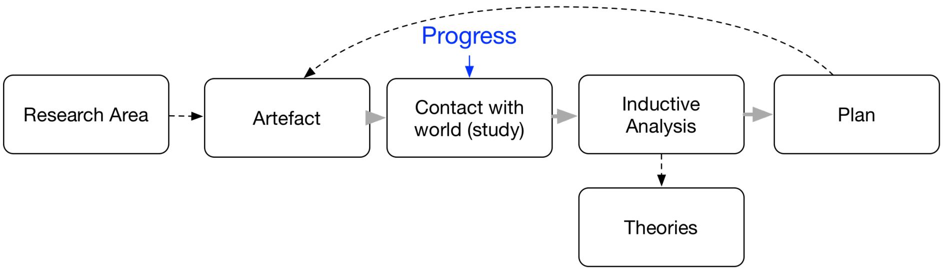

Figure 4.27: Current progress in the design cycle.

The studies for the MF system was in a similar format to that described in DTF in that they were typically informal discussions and workshops held throughout the design period. The system was once again assembled in the same space as the DTF system: a public space within Lancaster University. Again, given the interdisciplinary nature of the research carried out in the building by researchers and students housed there, the participants had a wide range of cultural backgrounds, ages, technical backgrounds and areas of expertise.

In addition to repeated informal trials and discussions throughout the design and build process, 12 mixed background participants were also asked to take part in using the tool in a more structured session. This session was divided into in three stages. The first stage consisted of a 10 minute open-ended exploration of the tool in which the user could get used to how the system worked. In the second stage each user was invited to complete a number of tasks of varying difficulty, and at each stage to describe aloud what they were doing and why. These tasks included: 1) select music files, 2) select only new music files, 3) select old music files and order them by rating. Each stage was repeated, first without the Filebots and secondly with the Filebots. Lastly the users were invited to take part in a discussion (first one-to-one with the researcher, then as a group) to voice their opinions on this tool, and more importantly this style of interaction. Whilst it took people a few minutes to grasp the mechanics of the Magnetic Tags, each user was able to successfully complete all tasks, and provide feedback.

One of the initial talking points discussed by each user was the way in which they became comfortable interacting with the projected files using Magnetic Tags. Many of those discussing the subject explained that this was due to the way that the files reacted to the tags in a way that mimics physical magnetism, and therefore they could predict reactions to certain interactions without learning the intricacies of the software. This is expected given the influence of TUI and NUI based concepts which are a foundation for this system. This positivity towards tangible systems agrees with themes taken forward from the DTF system discussed in the previous chapter.

The point that participants talked about with most enthusiasm (and spent the most time discussing) was the introduction of the Filebots, possibly due to the novel nature of the interaction. The discussions tended towards ideas of different physical techniques and their uses to represent various characteristics of the virtual world. For example, "should a hot file [Filebot] represent how new it is, or how highly rated the file is? Are some senses suited to certain types of information?". Similar questions were raised by everyone taking part in the individual discussions. In group discussion, the idea that different senses lend themselves to different types of information was discussed in more detail. Some people had preconceptions that certain senses should be used to represent certain information (such as temperature to be associated with relevance), whilst others believed that other mechanisms or senses were better suited (temperature should be linked to age, size should be associated to relevance). Interestingly, despite the fact that only two Filebots were in use during these sessions, every discussion demonstrated a good understanding of the Filebots, and a comprehension of what it was they trying to achieve. The discussion moved beyond what the Filebots could actually achieve, and into what they could achieve in theory. For example, the previous discussion regarding senses and file attribute included a mention of size: a property that the existing Filebots could not change. The Filebots were used more like props to provide context and a scenario to fuel discussion. In many ways it can be argued this more prop-like usage of the Filebots was positive as it allowed a more hypothetical, “what if?”, style of questioning and discussion which would have perhaps been limited with a fully working prototype system complete with hundreds of Filebots as per the initial idea. The use of props is actually very well used within the field of Design (Brandt and Grunnet, 2000, Brandt and Messeter, 2004, Buchenau and Suri, 2000), and indeed other fields as well. The “Wizard of Oz” experiment is commonplace in areas of experimental psychology (Dahlbäck et al., 1993) which allows the researcher to be involved in the interaction loop, simulating reactions from the system which would otherwise be too difficult or time consuming to do another way (e.g build some system to automate it). These techniques allow the exploration of ideas that perhaps are not realistic to achieve for the purposes of research (due to limited resources or technological advancements). Although the generation of props were not the direct intention of this work, the reaction that people had to the Filebots resembles that of interaction with props, and the discussions surrounding the Filebots were rich with idea generation. This concept of props is an idea that is carried forward into the next iteration of the design process.

To put the themes and discussion points focused around senses in context, the work of Physicist Tor Nörretrander can be considered.

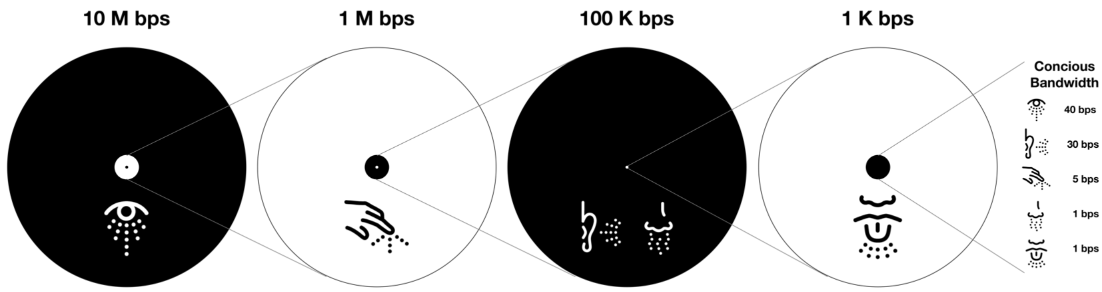

Figure 4.28: A Diagram showing Nörretrander’s sensory bitrates

Nörretrander explicitly introduces the different senses as components of an interaction, and discusses them much like digital sensors, complete with a typical bitrate[1] for each sensor. Figure 6.4 shows, from left to right: vision, touch, hearing, smell, and taste. Each of these senses is labelled with a bitrate, from 10Mbps down to 1kbps for vision and taste respectively. These values are based upon considerations of the number of sensory cells in the body for each

[1] A measure of the amount of digital information transmitted. It is typically measured in bits (1 or 0) per second

of the senses. Nörretranders further divides these rates for each senses into 'perceived' and 'sensed' rates which differentiate between what is consciously perceived versus that which is sensed. His reasoning for this is that only a very small percentage of what is picked up by the sensors (e.g sensed with the eyes) is being consciously considered by the person (i.e perceived by the brain).

If these concepts are related to the discussion points raised by those people interacting with the MF system, it is possible to see agreement between what was observed by participants and Nörretrander's work. As a result, it is possible to ask some questions about possible interaction scenarios. One of the first issues discussed by participants was the idea that at the start of the interaction - when all of the files were projected on to the table - too much information is being displayed at once, and it is hard to understand what is going on. According to Nörretranders work, the amount if information that is being presented to the user visually is perhaps surpassing the limit considered to be the conscious bandwidth of the sense of vision. The person likely still sees everything on the table, but it is not all consciously considered. Once the user places a Magnetic Tag on to the table, some of these files are removed and the level of information is reduced such that it is more easily understood. It follows that the level of information communicated has been reduced below the conscious bandwidth, and can now be more readily understood by the user. Throughout discussion the term "Information Bandwidth" was used as a term for this limit, and "Information Overload" as a term to describe the state when this limit is surpassed. It is important to note here that the term "Information Overload" in this context is not to suggest that there is too much information available to the user and as such the users ability to make decisions is negatively affected (as is the traditional use of the term (Gross, 1964, Speier et al., 1999, Toffler, 1971), which is synonymous with other terms such as "Infobesity" (Rogers et al., 2013)). Instead, it is used to describe a scenario where a user is unable to perceive any more information. The key difference between these two uses of the term is the consideration of what is "sensed" and what is "perceived" by a user.

Nörretrander's work highlights the possibility that this bandwidth is different for each sense. If this is the case, it is possible to ask questions such as "Can we alleviate Information Overload on one sense by moving some of the information to another sense?". The left hand side of Figure 4.29 shows the level of information transmitted in the case where visual Information Overload could be occurring; lots of information communicated via vision. This question allows the consideration of the possibility that Information Overload can be alleviated by changing the way that some of this information is presented, such as is shown on the right of Figure 4.29, and that Information Overload may occur on a per-sense basis. For this work the term "Bandwidth of the Senses" has been used.

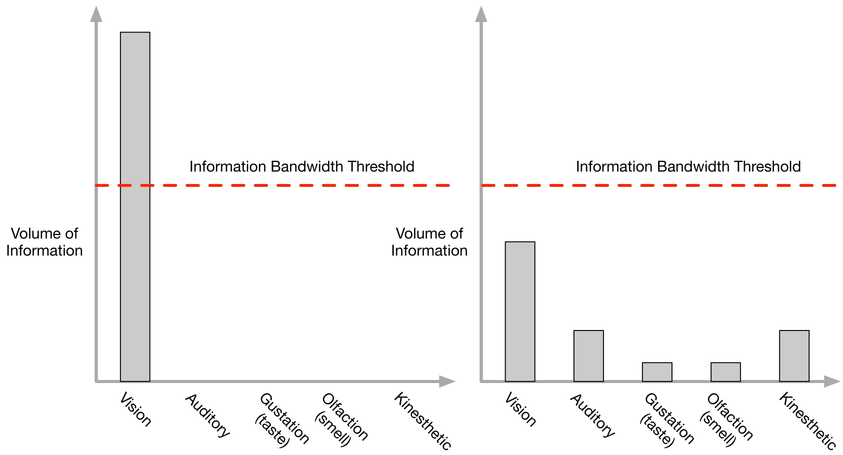

Figure 4.29: A Diagram showing Information Overload, and how it might be averted.

The differentiation between what is consciously perceived versus what is sensed is in some ways analogous to what Ishii and Ullmer (1997) describes as 'foreground' and 'background' media, with foreground being in the forefront of attention (perceived), whilst background is sensed but not consciously considered (sensed). In these discussions, these two ideas were grouped together under the name "Attention". Discussion often lead to comparison of different feedback mechanisms and senses and their effect on user attention. A loud, sharp noise can be considered attention demanding (or more foreground), whereas a subtle ambient noise could be considered less so (and therefore more background). In this noise example, sound can be both foreground and background depending on its format. The way that different feedback mechanisms affect different senses was discussed. A change in texture is, for example, more subtle than a loud noise. Users discussed ways that this difference in demanding attention between feedback mechanisms could be used when designing interactions, although only briefly.

In addition to the more generalisable concepts discussed above, a number of new interaction techniques were introduced by the study participants. The idea of "squashing & stroking" was introduced as a way for the user to quickly search for a file using senses other than vision. The user would position many Filebots close together and then stroke their hand across them assessing their physical attributes. If for example, heat was representing how recently the file was edited, and roughness represented the finished state of a file then the user would be able to quickly find the roughest and warmest Filebot. "Squashing" was discussed as one of may ways that the user could tell the system to perform some action on that particular file. A user could "squash" a Filebot to perform some compression algorithm on it, for example.

One of the people using the system was an avid music collector and expressed great excitement at the possibilities of managing their music using this sort of technique. They described the way that a physical interfaces like this could open up digital music collection to those that are less proficient with computers. He was particularly taken with the concept that his favourite tracks would be physically represented as somehow different from the others, making them easy to identify. Similarly, one idea that was discussed as a possible area for research was the investigation of how this interaction technique performs in some other real world scenarios. For example, exploring a Relational Database using "Magnetic Tags" as a query tool.

Moving on from the themes which are specific to the Magnetic Files study, the next section discusses the more generalisable themes taken on into the next cycle of the work.

4.5 Analysis & Themes

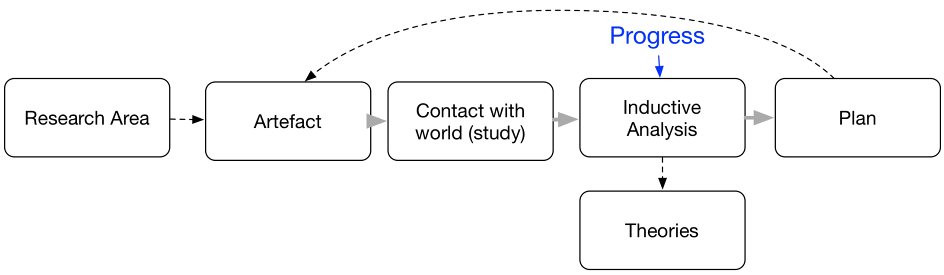

Figure 4.30: Current progress in the design cycle.

The concept of "Information Overload" re-occurred throughout discussion with many participants. It was mentioned that initially the tool looks intimidating as there are hundreds of files present on the work space - far too many for most people to actually process in any meaningful way. However, once the users placed magnetic tags onto the area, this effect was greatly reduced. Whilst the MF tool has one scenario in which this was definitely a problem (the initial state where lots of files were projected onto the table), further investigation into phenomena that cause Information Overload should be a part of the next iteration of this work.

All of the discussions at one time or another revolved around how different feedback mechanisms, and senses, were appropriate for different types of information. The inclusion of some element that allows this to be further explored should be part of the next iteration of this work.

Some of the discussions indicated that a state of Information Overload can be reduced by representing information using a different sense. This suggests that the Information Bandwidth - the maximum amount of information perceived by a user - operates on a per sense basis.

Nörretrander and Ishii both discuss attention, but from different viewpoints. Whilst Nörretrander focuses on the senses and what can be consciously perceived, Ishii discusses that different types of media may lend itself to certain levels of interaction (be more or less attention demanding). This differentiation is highlighted by some users through discussions: some suggested that certain mechanisms are more suited to a (to use Ishii's terminology) background role in the user's attention, whilst others are more foreground in nature. This concept of attention should be considered in the next iteration of this work. What feedback mechanisms and senses suit foreground, and which suit background? Can transition between these two states be achieved?

Finally, the more prop-like usage of the Filebots highlighted the benefits of such an approach: props give context and encourage discussion with similar framing. The ideas and discussions generated throughout this chapter were arguably richer because the discussions moved beyond what the technology allowed and onto what could be imagined. Given the benefits of this approach, the use of props to encourage discussion should be considered an important aspect to bring forward into the next iteration of the design cycle.

Chapter 5

ANTUS

5.1 Paper

This chapter was originally written as a paper for academic consumption and presented in Brighton, UK at Design Research Society (DRS) [1] (Gullick and Coulton, 2016). It was generally well received, and in addition to being a thoroughly enjoyable and informative experience, I hope that some of the feedback and clarification needed throughout the review and discussion process has positively impacted this chapter.

5.2 Planning

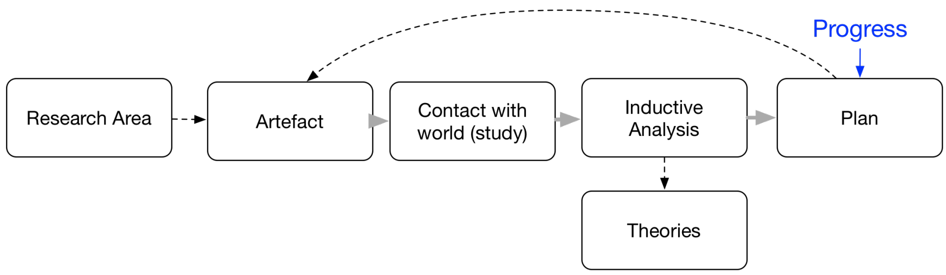

Figure 5.1: Current progress in the design cycle.

[1] http://www.drs2016.org/

Table 5.1: A table outlining the themes Antus aims to explore, how these themes could be investigated.

Theme	How Investigated?
Different senses	Mechanisms based upon multiple senses within Interaction.
Information representation	Different types of information available. e.g. Quantatitive and Qualitative.
Information Overload	Lots of Information.
Attention	Information with different immediacy needs.

As the iterative research process outlined in Chapter 2 reached its third iteration, the emergent themes that this work endeavoured to identify began to consolidate. The work carried out during the Digital Terra Firma (DTF) project highlighted the natural expectation of users for objects to consist of more than just visual elements. The MagneticFiles (MF) project brought additional senses into the interaction, and, in doing so, it highlighted a number of avenues worthy of further consideration: the different feedback mechanisms and on what sense they are based; the types of information that different senses may lend themselves to; the causes of Information Overload, and ways to avoid it; and issues of how attention is considered during this interaction.

The project described throughout this chapter further investigated these key areas, incrementing upon the work of previous chapters. It drew upon the success of the use of props experienced throughout the MF project and aimed to build an artefact with the goal of encouraging discussion focused around these points.

In a similar fashion to the DTF and MF, the Antus project began with early brainstorming and work-shopping sessions. The goals of these sessions were to highlight key areas upon which this project needed to focus. Given the themes brought forward from the previous iterations of the design cycle, Table 5.1 outlines potential requirements for this system that can be concluded.

Throughout the initial work-shopping session, many ideas were discussed for the theme of this project and the effectiveness at matching the above requirements were compared. These spanned from concepts of landscape management closely mirroring the system discussed in DTF, to Virtual Reality (VR) systems, to various interactive games. Eventually, the theme of an augmented tabletop game was considered the most appropriate. In a game, different types of information about game characters can be represented: qualitative and quantitative; information that needs taking into account immediately; and information that does not. A game environment affords flexibility to tailor these aspects. At the same time, games offer an easy and engaging entry point for people participating.

5.3 Artefact Creation

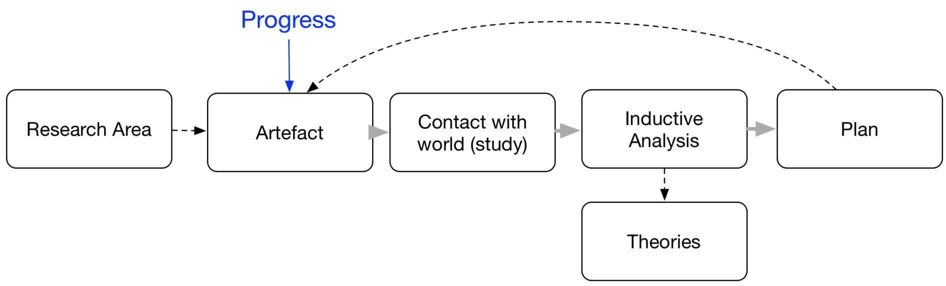

Figure 5.2: Current progress in the design cycle.

Much of the software and hardware needed for the successful creation of augmented tabletop games had already been constructed throughout the DTF and MF projects. A game around rivalling ant colonies was chosen as the basis for this work as it appeared to offer the greatest variety of possible interaction scenarios and was most positively anticipated by those taking part in the planning session. A game centred around rivalling ant colonies could closely meet those requirements outlined above as, by involving ants, it can have a very large number of entities controlled by the computer (each is known as an Non-Player Character (NPC)). With such large numbers of characters, large amounts of information can be displayed to represent and accompany them.

As mentioned above many of the elements from the previous two projects were incorporated: the projector; Microsoft Kinect and Red, Green, Blue: a system for representing the colors to be used in an image (RGB) camera; Augmented Reality (AR) markers; the second version of the chassis from the Filebot; and relevant software to synchronise the physical and digital spaces as much as possible. Antus can considered as a form of augmented tabletop game. Although it is not the focus of this research, there are unanswered questions regarding tabletop games that this work may help to answer such as: “How should information be presented in AR tabletop games?” Augmented tabletop games have been the subject of much research in recent years (Kojima et al., 2006, Leitner et al., 2009, Magerkurth et al., 2004) although the majority have been used to highlight novel technological interactions and they have not considered the information objects within the game may be required to represent. Bakker et al. (2007) identified that players generally preferred physical objects over virtual ones, and this reflects the feedback throughout the research of the DTF and MF systems. However, the issue highlighted by Magerkurth et al. (2004) of understanding whether feedback should be physical or digital within the context of augmented tabletop games remains unanswered.

Antus has been designed in such a way that it has information that could be characterised as "hot" and considered to be available in the players "foreground" of attention, and information that could be characterised as "cool" and suited to display in the user's periphery (or "background"). This data, and its representations is looked at more closely in Section 5.3.7. As discussed previously, the aim was to use this game as a design stimulus for a workshop focused on representing information in the game space in a physical way.

5.3.1 Primary Game Mechanics in Antus

Essentially, Antus was designed as a "God Game", in that the player controls the game on a large scale, as if they are an entity with divine or supernatural powers. Each player takes control over a colony of ants, and is charged with controlling the actions of that colony. The main goal of the game is to provide food for the Queen ant and her nest via the use of Farming ants which can forage for food, whilst simultaneously hindering the progress of rival colonies with the use of Soldier ants and the manipulation of the physical surface.

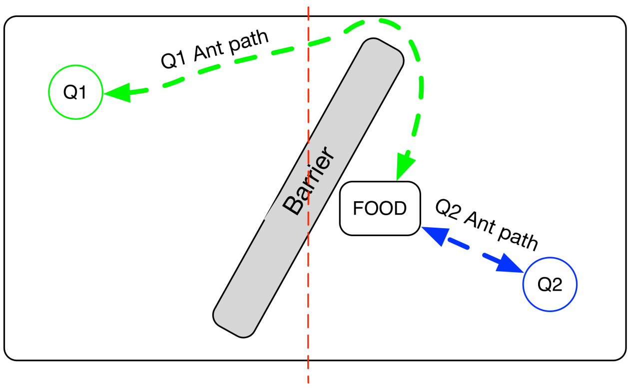

Figure 5.3: An diagram showing the path of ants in Antus.

Figure 5.3 shows an illustration of these primary game mechanics. The two nests, or Queen ants are represented by the circles labelled Q1 and Q2, for player one and player two respectively. The path of ants (shown as dashed lines of similar colour to the Queen of each player) from each nest towards food can be seen to avoid physical barriers on the game space. In this scenario, player 2 has the advantage as their ants have a shorter path to get food and return it to their Queen. Food acts like a currency in this game, and it is used to create more ants and feed the Queen. The following paragraphs describe the primary game mechanics outlined in these early sessions.

As the players compete, they can add, move, or remove physical objects to adjust the barriers which the ants must avoid. They can also place objects representing food down onto the tabletop, which the ants can consume (and therefore these food deposits deplete over time). In early tests, players would exploit a loophole in these rules by placing food directly on top of their own nest, which limited inter-colony interaction. To combat this, the idea of a territory divide - shown by the red line - was introduced. Players can only place food in the territory closest to their opponent.

Ants in this game can exist as one of the following three forms: the Queen ant represents the nest and must be fed at all times; the Farmer ant which can collect food; and the Soldier ant which can kill ants from the opposing colony if they cross paths (and bring them back as food) but cannot collect normal food. The tactics of deploying the Farmer and Soldier when playing Antus is in addition to the physical manipulation of the game space and the strategic placement of food and players must also manage the production of Farmer and Soldier ants.

The aim of the Antus is to last longer than your opponent. As the food resource is constantly used for keeping the Queen fed, and creating and feeding other ants, when this resource runs out for either player they lose and their opponent is deemed the winner.

5.3.2 Construction

So far the ideas and theories behind Antus have been discussed, but the construction has not. This section addresses this, and is split into Hardware, and Software sections.

5.3.3 Hardware

Fortunately, much of the hardware used throughout Antus was able to be re-purposed from the previous projects discussed in this thesis. The same configuration of Microsoft Kinect, Projector and AR markers was used as throughout the DTF and MF projects, and the table used for the MF project was used as the flat space upon which the game was played.

AR Markers were used for two key purposes within Antus. Firstly, they were used as objects to position food onto the table. These food objects acted as a reconfigurable tool (discussed in Section 3.3.1) and painted food onto the table. Once the food had been painted, the object and corresponding marker could be removed.

The second use of AR markers within Antus was as Meta Parameters discussed in Section 3.3.1 of the DTF project. Certain markers could be placed onto the table to issue commands to the ants. For each player, AR markers existed to allow them to start/stop creating ants, and prioritise Farmer/Soldier ant production.

The Queen ant throughout this work was represented physically, unlike the Farmer and Soldier ants which were simply projected onto the space. The Queen ant for each player was represented by a model ant based upon the second Filebot chassis discussed in Chapter 4. It was decided that the Queen should be represented physically as it offers opportunity for feedback mechanisms other than the visual, which was a requirement as indicated by Table 5.1.

5.3.4 Software

Whereas the hardware for Antus was mostly re-purposed elements from the previous projects discussed in this thesis, the same was not true for the software. Having said this, most of the low-level software that interfaced with the projector and Kinect to synchronise virtual and physical environments was re-used, but a number of additional modules were needed for this project.

5.3.5 Game Engine

The previously discussed projects used JMonkeyEngine as the engine behind what is shown via the projector. JMonkeyEngine is primarily written for use as 3D software. When trialling early versions of Antus it became apparent that the Three Dimensional (3D) nature of JMonkeyEngine was was detrimental to the speed of the game. With so many NPCs displayed at once, the extra dimension of calculation was causing each frame of the game to take too long to calculate, resulting in jittery and frustrating gameplay.

To alleviate this issue, an extra module was added to the software. This module performs the same job as JMonkeyEngine did in the previous projects but it performs this task in only in two dimensions. This module is based upon the core library of Processing[1]. As such, it allows low level access to portions of the graphics pipeline (the sequence of steps used to create a 2D raster representation of a 3D scene), such as Open Graphics Library (OpenGL) textures (a commonly used, fast, but arguably non user-friendly graphics Application Programming Interface (API)). With this switch, the game ran without issues at 30 frames per second, resulting in the intended game-play experience. Figure 5.4 shows how the software has been augmented since its advent in the DTF project (see Figure 3.10 for the original).

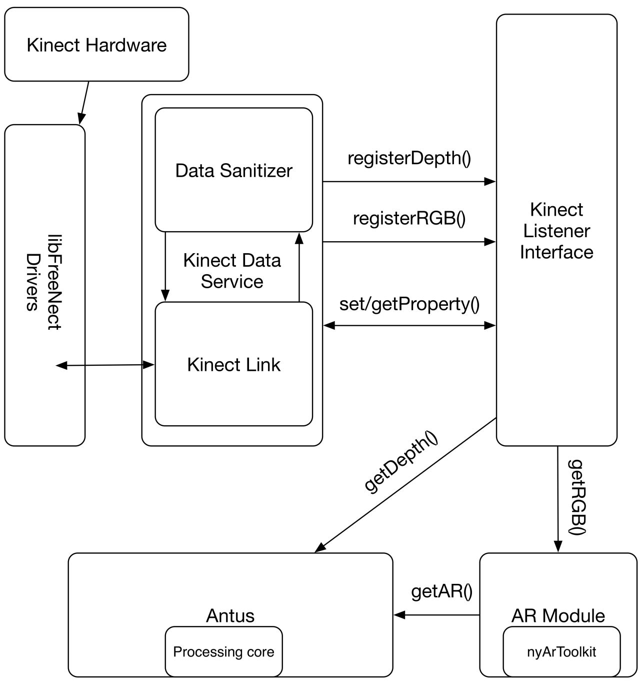

Figure 5.4: A diagram showing the flow of information through Antus.

[1]https://processing.org/

5.3.6 Ants

Ants in Antus are designed much like agents in agent-based modelling: a class of computational models for simulating the actions and interactions of autonomous agents to assess their effect on an overall system. As such, each Ant was programmed to mimic the behaviour of a real ant. For example, they will die if left for long periods unable to find food.

Real ants rely on the use of pheromones to navigate. Simply, if an ant find something good (like food) it begins leaving a trail of a positive pheromone, whilst it leaves a negative pheromone if it comes across something bad (like danger). Ants that are new to the region will follow good pheromones (the direction to follow can be determined by strength of the pheromone in each direction), and avoid bad ones.

Ants in Antus mimic this behaviour by leaving trails that inform other ants. Rather than good and bad pheromones, Antus ants leave "food" and "home" pheromones (called trails). Those ants looking for food will follow food trails left by other ants which have just found food, those with food will follow home trails .

At any point in time ants in Antus will make a decision about where to go next. This decision is weighted by nearby food, pheromone trails, elevation, and the presence of enemy ants. In addition to this, each decision has a small chance to ignore logic and act at random. The resulting paths taken by the ants throughout game play begin by looking random, but quickly organise themselves as the ants communicate with these trails.

5.3.7 Running Antus

Up until this point, Antus has only been discussed in theory, or the way that it was built or programmed. This section describes how these elements came together to make the game, with images to illustrate.

5.3.7.1 The Ants Characters

As discussed, ants can exists as one of three forms. A Queen ant is represented physically with the use of the Zumo robots (used as the second version of the chassis in MF). These Queen were largely stationary in the game, representing the nest of the ants. In early versions of the game the Queen ants were able to move if they were not receiving enough food from their colony. Despite this fairly limited role, the Queen ants became the centre of much discussion (see Section 5.4). The remaining two types of ants in the game were represented virtually through projection on to the surface. As so many characters were computed at one time, these ants were represented using a simple rectangle. When carrying food (discussed in Section 5.3.7.4) these ants were accompanied by a smaller rectangle. These can be seen in Figure 5.5.

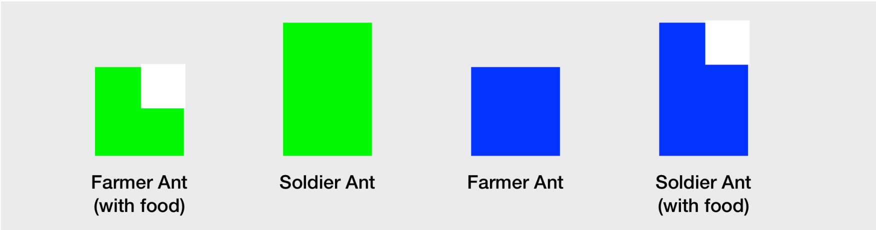

Figure 5.5: Antus Farmer and Soldier Ants.

5.3.7.2 Obstacles

Ants in Antus were programmed to be lazy: they will not travel uphill, and will always navigate around objects instead of going over them. Using the Microsoft Kinect, a map outlining the height of all objects on the tabletop was built, and this map was used to inform each ant's decision on where to move next. Therefore, as a player, it is easy to manipulate the path of the ants (both yours or your opponents); you can simply place an object in their way and force them around it. Due to the internal workings of the Microsoft Kinect, any object that reflects Infra Red (IR) light works for this, and as such players tended to use whatever was at hand for this purpose (e.g. rolls of tape or books).

5.3.7.3 AR Markers

AR markers were used for two purposes within Antus. Firstly, they were used to specify where a player would like to place food on the table, and secondly they indicated what actions a player would like their colony to perform. The various options are discussed in Section 5.3.3. These AR markers were simply printed onto paper, and reacted to when they were within view of the camera in the Kinect.

5.3.7.4 Food

Food in Antus was placed onto the workspace via the use of appropriate AR markers. Once these markers were detected (and stationary so they could be moved into position without leaving a trail of food) the system would mark the area with a deposit of food. Food was represented by a white area. As ants reach the food, they removed some of the food, thus reducing the size of the deposit. This can be seen in Figure 5.6 which shows some of the Antus gameplay.

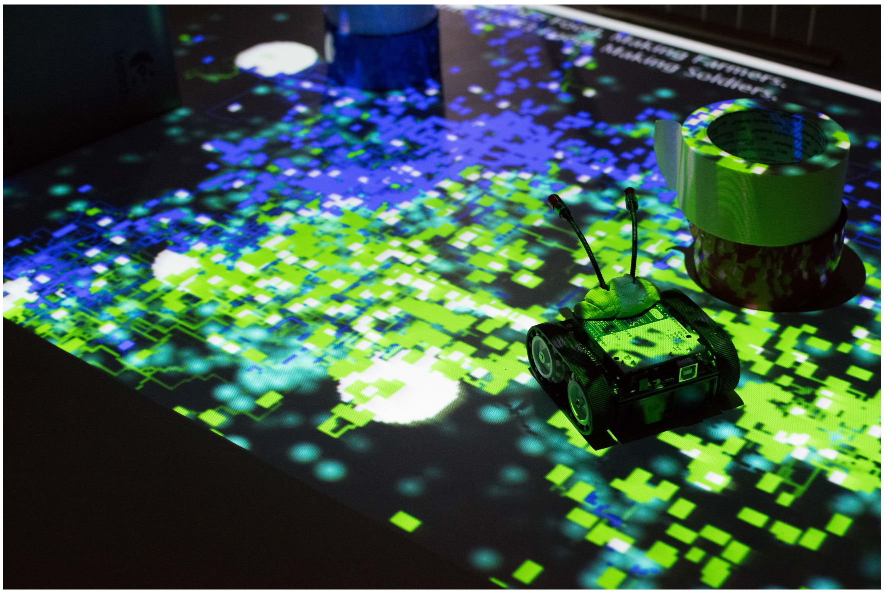

Figure 5.6: Antus gameplay showing food. The two colonies (green and blue) can be seen alongisde the Green team Queen ant, some obstacles and food deposits.

5.3.7.5 Pheromones

As discussed in Section 5.3.6, the Soldier and Farmer ants are programmed in a way that they act like autonomous agents. Whilst the player has control over the general direction or tasks that their colony performs as a whole, they do not control individual ants. These ants navigate via the use pheromones, which are represented by projected trails left behind the ants. Over time the trail fades, and so the complex map of trails left by ants changes over time as the food locations, obstacles, and ants change. Whilst these pheromone trails are visible in the game (such as in Figure 5.6) they are most clearly visible in early versions of the software as shown in Figure 5.7.

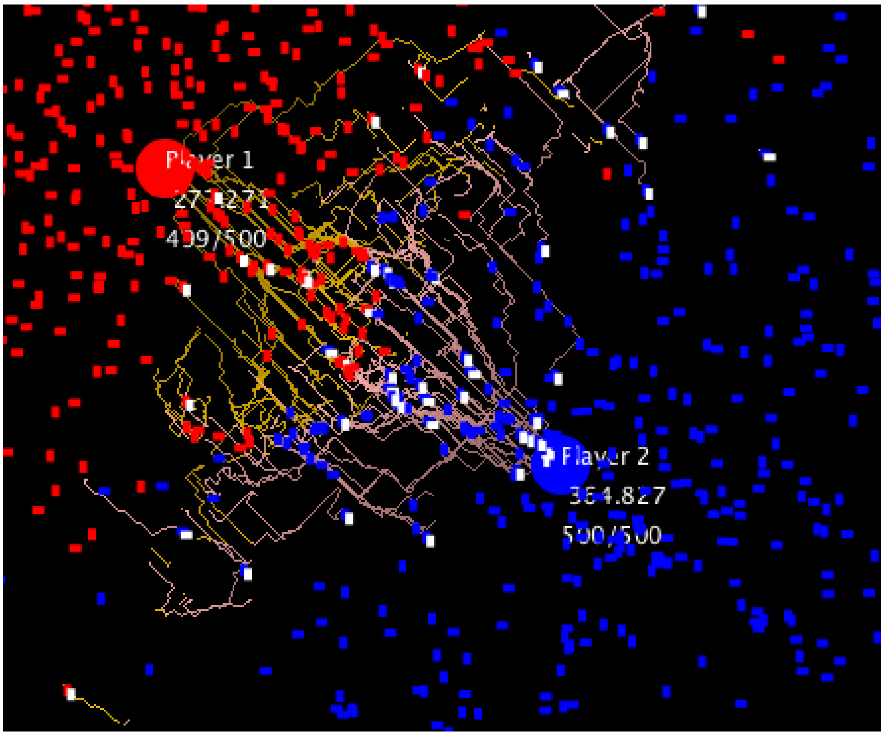

Figure 5.7: Early Antus gameplay showing pheromone trails.

5.3.8 In Game Information

The level of information expressed to the players throughout gameplay is specifically intended to be very high, so as to encourage Information Overload. Figure 5.8 shows some of the explicit information visually expressed to players throughout gameplay. The text shows which team is which (with coloured ant icons), how many ants, how much food, and the current command of each colony.

Figure 5.8: Explicit information available in Antus.

All of the previously discussed phenomena are represented visually. In a typical game, over 2000 NPC ants which may be Soldiers or Farmers, and may or may not be carrying food are displayed. Additionally, pheromone trails are left for each of these ants, food locations are shown, as well as locations where ants have died (these can be seen as the cyan coloured splats in Figure 5.6).

In addition to these visual elements, a few different sounds are also part of Antus. Each time a new ant is created, food is collected, or an ant dies, a sound is played to represent each scenario uniquely. Much like MacKay (1999) discusses how the sounds of the room were used to gauge the level of work and tension in the room of air traffic controllers (discussed in Chapter 1), these noises allowed the users to assess the state of the game. In actual use, so many noises were being played at once, they produced more of an ambient tone that represented the state of the game rather than individual noises to represent events.

5.4 Contact with the World

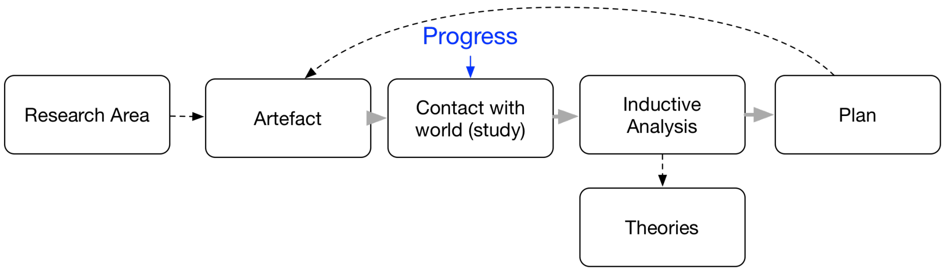

Figure 5.9: Current progress in the design cycle.

It would be possible to address the previously highlighted questions relating to the senses, attention and information by attempting to build different solutions to problems and then testing them with players. However, given the success with the prop-like interaction within the MF study, it was decided that a participatory design approach would allow a wider range of options to be considered and would facilitate conversation with players about interaction within hybrid physical/digital game spaces.

For the participatory design workshop, Antus was set up in the same space as DTF and MF, inviting participation from a set of students and staff at Lancaster University with diverse academic, professional and cultural backgrounds. Each participant was invited to play Antus as it is described above. Initially this gameplay was one participant against a researcher who outlined the rules of the game, and later the gameplay was between participants without input from the researcher. An image showing some such gameplay between participants can be seen in Figure 5.10.

Figure 5.10: Antus Players.

The players were then invited to comment on the current in-game information and then to consider alternate ways of providing that information. Whilst many people played (and enjoyed) Antus due to the public space, only eight participants (six male and two female) were available to take part in the extended workshop event (gameplay, discussion, and modelling).

Figure 5.11: Antus Workshop.

Physical prototypes of alternate ways of representing this information along with suggestion of new information that could improve overall gameplay were encouraged. Whilst systems have been created that allow prototyping of physical game objects (Marco et al., 2012) these were aimed at games designers and offer a limited range of ways in which to represent information. Therefore it was decided that providing players with a range of craft materials would allow them to express their ideas much more freely in the given time (Hare et al., 2009). Some of these craft tools can be seen in Figure 5.11. A sample selection of some of the prototypes produced is shown in Figure 5.12 (general feedback prototypes) and Figure 5.13 (Queen ant based prototypes). This workshop offered many insights into how players approach the problem of physical data representation, and gave people the opportunity to explain some of the less obvious design decisions. The most interesting and relevant insights are as follows.

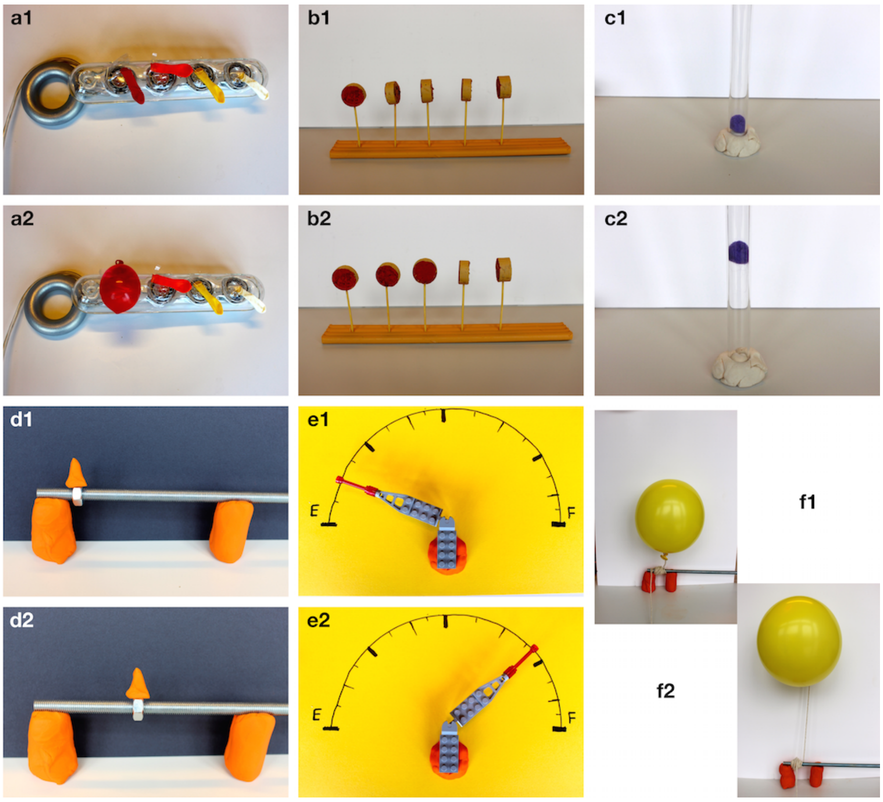

Figure 5.12: Feedback mechanisms designed by participants.

Figure 5.13: Queen ant feedback mechanisms.

- Users expected the robot actors to have emotion and this emotional state has been shown to be important aspect to gameplay (Barakova and Lourens, 2010). A number of the prototypes built by participants support this result. From the designs generated in the workshop, texture and sound are often related to a state of emotion when used with robotic actors. Most participants chose to use rough textures, such as 5.13b1 and 5.13b2, and fast noises to represent a negative mood and smooth textures and slow noises to represent a positive mood. Whilst this could be a result of the insect based theme of the game, it nonetheless supports the notion that robotic actors have emotional states.

- The relative difference of information is often more important to users than specific value of the data. Many of the prototypes were designed in such a way that it represented relative concepts such as “more than my opponent” or “doing well” rather than to represent specific values to the users. When questioned about this, one participant explained: “the amount isnt important, its being able to easily see your relative position to your opponent that is important”.

- “Glanceable” feedback was important to a number of participants as they wanted to spend more time considering their strategy and playing the game rather than exerting effort to decode the information. Additionally, many of the feedback mechanisms were designed to operate in the background: players did not want to be interrupted to be told the state of the game, and instead wanted to choose when to get feedback by looking, touching, or accessing feedback which was more ambient so that they can get a sense for the state of the game. The creator of prototype 5.125b said they wanted to recreate “those mechanical displays you used to see in train stations or bus stations” because “you dont have to keep watching them as the noise tells you when something has changed”. This noise is effectively ambient information with an alert that helps bring the information from background to foreground.

- Multiple senses can be used simultaneously to perceive information. For example, one participant designed a feedback system that utilised a speaker, Light Emitting Diode (LED) lights, and an inflating balloon to represent different aspects of the game within one feedback device (ant death or new event, food low warning, inflation to represent the level of food respectively). The creator explained it was easier to understand than a purely visual feedback device. Often one sense was used as a cue to let the user know that new information was available. In this case, although the balloon inflated to represent the level of food the colony had collected, the LED light indicated that a certain low threshold had been reached, and that the user needed to pay attention to the balloon.

- A scale on a feedback device is not always necessary as some participants chose to not include a scale, and just to represent a state change and a relative direction.

5.5 Analysis & Themes

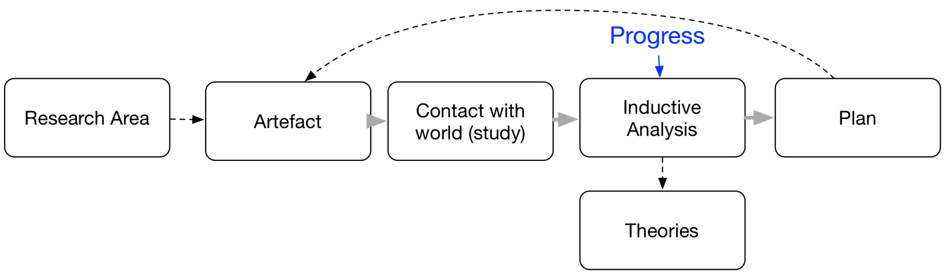

Figure 5.14: Current progress in the design cycle.

One of the most interesting insights that builds upon previous themes was the way that people managed their attention throughout this process. Attention was treated much like a resource, and the prototype feedback mechanisms produced in the workshop by players were designed in a way that limits the amount of attention required. Players designed "Glanceable" feedback mechanisms that allowed the choice of when to receive new information. Other mechanisms would use one technique to indicate to the player that new information was available, and the player could then choose when to investigate the new information. Even then, when information was presented to the users, the form of this information was changed to limit the level of information presented. Explicit values and scales were removed, and relative information was displayed. This indicates that the players are already quite fluent in treating their attention as a resource. The feedback mechanisms that achieved this spanned multiple senses, at times using one sense (such as a sound) to indicate that other information was available via another sense.

However, it was noticeable that players were focused on what could be more easily seen, touched, or heard as the main channels for feedback, although in later discussions additional senses, such as smell, were described as possibilities to indicate certain types of information.

This may be due to the physical crafting nature of the design workshop: it is hard to represent something as abstract as a smell with a physical prototype.

Whilst further investigation is needed to fully understand the effect of the senses on attention, the insights discussed above offer a deeper look into the way that people choose to interact in this space, particularly the way in which they wish to receive information, and the amount of attention it requires. These insights are incorporated into the collection of themes, and resulting framework in the following chapter.

Chapter 6

Reflections and conclusions

6.1 A Review

The premise for the work carried out throughout this thesis is the continued development and evolution of the interaction that occurs between humans and computers. One of the key assumptions made is that technology is improving and evolving at a rapid pace, and as it does so too must our practises for interacting with it. Until relatively recently, what had existed as two clearly distinguishable and conceptually separate spaces of the physical and the virtual are now becoming more difficult to distinguish from one another. Traditionally, the separation of these two spaces by a screen of some sort (consider the past 20 years of video game consoles attached to screen for gameplay) has made it relatively easy to distinguish between them. As technology has improved and become more ubiquitous in nature, this easily defined distinction between what is real and what is virtual is no longer so simple. Whist technologies that have enabled some elements of this spatial duality have existed for a number of years (e.g Virtual Reality headsets), it is only recent that these technologies have become truly mainstream and sufficiently technologically advanced to start exhibiting the level of duality that is discussed in this work. Sticking with games consoles as a solid example of this, the latest generation of popular gaming consoles, such as the PlayStation 4, Xbox one, or Nintendo Wii, all offer some sort of interface for translating physical movement into the digital (the Playstation eye, Xbox Kinect and the Wiimote respectively). Similarly, VR headsets are now much more common : hardware such as the Oculus Rift, the HTC Vive and even technologies such as Google

Cardboard[1] allow easy access to VR and AR experiences, with major graphics companies such as Nvidia acknowledging and supporting VR as an important feature. Even without specialist hardware, recent mobile applications such as Pokemon Go, or game objects like Activision's Skylanders are allowing people to experience hybrid spaces where the physical and virtual are no longer necessarily separated by a screen.

With this hybrid nature comes the difficulty of understanding and designing interactions. From the interactions carried out throughout this work, and my experience in general, people often decide on the physical/digital nature of an object before interacting with it, using these preconceived ideas to dictate how they understand the object and therefore which actions they believe can be performed with it. As this line between these spaces becomes less clearly defined, it follows that so too must the rules for interaction.

The problem explored throughout this thesis is arguably one of mental problem framing: "how do we understand these hybrid spaces" and "how do we design for interaction within this space" are questions which address our preconceived notions about virtual and physical spaces. Conceptual framing for any problem is critical, especially in these more emergent areas: the conceptual framing will determine what questions are asked, and how the answers to this questions are obtained (as discussed by Schön (1987). This focus on framing is reflected throughout this thesis, and and has been explicitly examined and questioned: framing is, at its core, influenced by an epistemological positioning that plays an important role in decisions regarding methodology, methods, and how new knowledge is interpreted and used. Being explicit at every stage from epistemological standing to artefact creation to themes means that this work can more effectively bring the research community closer to a proper framework for working and designing interaction within these hybrid physical/digital spaces, especially in scenarios where information is being passed within the space.

In this chapter each of the themes carried forward throughout each iteration of the research cycle (each project chapter) will be discussed in detail, and then drawn together into a framework. This framework focuses on the senses, and the way that information is conveyed through sensory channels to the user (how it is sensed, how it is perceived, and how the user is expected to react). Whilst additional work is needed to further formalise these theories and concepts into a more solid framework, this work is presented as a step towards this goal. This

[1] A downloadable cardboard headset which acts as a Virtual Reality Headset with a compatible phone. See https://vr.google.com/cardboard/

thesis concludes by reflecting upon this framework, the research process, and the experiences gained through the undertaking of this research with a brief outline of what would be considered appropriate future research pathways.

Looking back at the DTF project discussed in Chapter 3, the DTF system is presented as an attempt to "synchronise" digital and virtual spaces with the goal of creating a hybrid space in a relatively controlled environment. The project was a success in many ways and it provided much of the planning and research into technology, and techniques that would form a foundation for projects later in the research. Through interaction with various people throughout the course of the project, a number of interesting phenomena were noted: the impressive ability for tangible systems to work in a cooperative environment; the speed at which people could learn and adapt when interacting with the physical; and the way that the chosen technology works, such as AR markers. Alongside these interesting insights, arguably one of the most important outcomes was the highlighting of a failing of the DTF system: the synchronisation of physical and virtual spaces across multiple senses. In hindsight, this is not surprising considering the technology that is employed to synchronise these two spaces: the Microsoft Kinect, and a projector system, are both technologies which are designed with a heavy focus on visual aspects. This shortcoming of the DTF system was discussed multiple times during the design process: some people commented on the way that they could place virtual 'fire' using the available tools, yet this fire did not emit any heat, others mentioned the way that the system was focused purely on the way objects look and didn't represent well any other properties. Whilst it can be argued this bias towards what is visual is a result of technological availability, techniques do exist that could be used to investigate the other senses (some of these techniques are investigated in part in the MF project).

The concept of utilising the different senses is interesting and is a key factor for this work. Figure 6.1 shows how different points along the Virtuality Continuum (a spectrum ranging from the entirely virtual to the entirely real) currently differ in their sensory associations. Each side of the continuum has been illustrated with representations of how we know to interact with that space, yet a region of overlap of the physical and virtual exists.

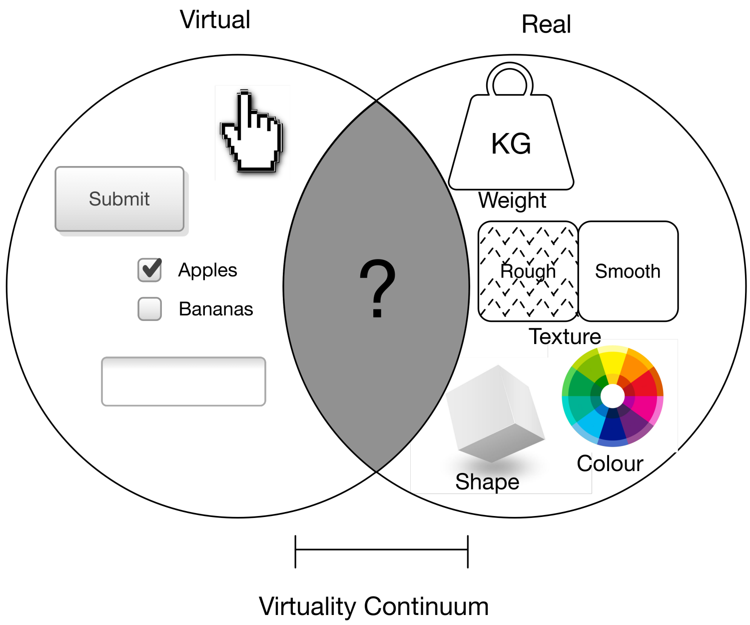

Figure 6.1: A Diagram showing the different spaces and how we understand them

If we look towards the right of the diagram, inside the circle representing the physical, Figure 6.1 depicts weight, texture, shape and colour as some of the ways that users may assess the actionable properties of a physical object. These others are not included in the diagram for reasons of space but include properties such as rigidity, temperature, density, and size.

Towards the left of the diagram is the purely 'virtual space'. The way this space is currently understood appears to somewhat reflects that of the physical. In Graphical User Interface (GUI) based applications, we create virtual buttons by drawing a 'shadow' around an object indicating that it has a 3D button-like shape (so is able to be pressed) and that it has two states: pressed, and not pressed. The same appropriation of the affordances (discussed in Chapter 1) used in the physical world are apparent in many aspects of virtual interactions. Most of the standard widgets often seen in GUI environments also use shadowing to give a 3D appearance and indicate the option of interaction. This appropriation can be seen as the result of an attempt to bridge the disconnect or conceptual gap that exists between the binary world of computers (virtual) and people (physical).

The 'Do? Know? Feel?' cycle, outlined in Verplank's Interaction Design sketchbook, is a a way of conceptualising this process. This process is discussed in more detail in Section 1.

This represents the ways in which users understand how to interact with the objects around them. This cycle is built upon the concept of affordances: the way in which our experience of the physical world, and our cultural background indicate to us the actionable properties of an object. Moreover, the "do?know?feel?" cycle also subsumes concepts of learning and experience, indicated here by the Map and Path icons (an idea of mental map building first introduces by Lynch in the 1960s (Lynch, 1960)). Finally, the concept that different types of object have different actionable properties are indicated by Verplank's labels of 'handle' and 'button', which differentiate continuous vs discrete media, and 'hot' and 'cool' media based upon the work of McLuhan (1994) which differentiates between sensing and knowing. This is discussed in more detail in Chapter 1. It is evident that there is a wealth of knowledge that describes the way people understand how to interact with objects in the physical world: it is largely influenced by their experiences of the laws of physics, and the fundamental properties of the world in which they exist.

Having said this, additional ways have developed over time to understand and interact with virtual systems and a number of successful metaphors have stuck within the virtual world. Although over time many different metaphors have existed, the ones which have persisted, usually mimic aspects of real life: the computer mouse allows the user to point and select using a virtual hand for example, a metaphor which was widely accepted at a time when touch screens were not common. Similarly, the widely accepted 'Desktop Metaphor' provided a logical way for people to organise their virtual 'files' into 'folders' on a 'desktop' much as they would in real life, on a physical desk (although this metaphor did not transfer well to mobile devices). Most successful interaction metaphors in some way mimic processes or attributes from the physical world. This is logical as people will have certain preconceptions about the actionable possibilities and reactions that objects have in the physical world. If these actions and preconceptions are mirrored in the virtual world then the person taking part in the interaction would be more likely to already have at least a basic understanding of the 'rules' of the system; the effects of their actions, the state of certain elements, amongst others. It can therefore be expected that they will feel more comfortable with this interaction technique than an alternative that doesn't offer familiarity. It is worth noting that much of what is discussed here is a reflection of what happened 'on-screen'. This isn't a result of overlooking alternatives to screen-based digital interactions, but is instead a reflection on the reality that most of our interactions with the digital realm occur on a screen of some sort.

However, as reflected by the discussion throughout the DTF project: if a physical object is assessed based solely upon its visual aspects opportunities to transmit a wider range of information using the other senses could be lost. Much of our knowledge is subconsciously acquired through interaction using all of the senses in a semi-exploratory manner: for example a user can feel that they are pressing a button too hard, or that a rotary dial only rotates so far, and at certain intervals. If these type of controls are to be represented digitally, all of this information must be represented using some other means (usually visually), or is not communicated to the user.

As this work focuses on hybrid digital/physical spaces it follows that the division of the senses into each respective space (see Figure 6.1) should be carefully considered in the creation of a hybrid space. A combination of senses hold great importance to this hybrid space, and this is reflected in their inclusion in the second iteration of the research cycle: the MF project. Whilst this second artefact focuses on file system interaction rather than virtual world design, it still builds upon these same interesting concepts as the DTF project, using much of the DTF system as its foundation. Whereas DTF focused predominately on the visual, the MF project expands upon this by introducing techniques that allow exploration of the senses through the use of multi-sensory “Filebots”.

Whilst the MF project is based upon filesystem exploration, this context is only important as it supplies a scenario in which a wealth of virtual information is available, information that exists in many forms including file type, file size, age, location, state and word count. Whilst the project was limited by time, money, and technology in terms of the number of “Filebots” that could be created, it does offer a number of insights into this hybrid physical/digital space. A number of these insights are tightly coupled to the MF system itself: the concepts of stroking and squeezing as forms of interaction, or the discussed interaction scenarios in which music or photos can be easily navigated.

However, there are a number of additional insights which seem more generalisable in nature. The discussion carried out throughout the project cycle indicated that there is an upper limit to the amount of information that can be conveyed to the user using each particular sense: each sense has a so-called bandwidth, which, throughout the entirety of this work, has been termed “Information Bandwidth”. When the maximum bandwidth is surpassed, the level of information presented to the user is too great to be usable: a state referred to in this work as “Information overload”. Whilst this may be obvious for channels such as the visual (lots and lots of visual information may be hard to interpret), it does bring in to question the bandwidth of the other senses too: how much information can be conveyed using something like smell compared to vision? Whilst is is widely accepted that our senses are prioritised differently in the brain, and have different number of sensor cells for each, the most easily transferable way of conceptualising this is presented by physicist Tor Nörretrander. Nörretrander presents the senses of a human body as typical computer sensors, complete with a bitrate of information. This is discussed in detail in Chapter 4. Treating these senses like typical computer sensors leads to a number of interesting questions: 'what is the maximum bitrate?', 'what happens if we go over this bitrate?' and 'what type of data can be transmitted/received using each "sensor"?'.

Also discussed in Chapter 4 is the way that Nörretrander, alongside other academics, consider attention to work in this interaction. In essence, attention is divided into foreground (that which is consciously perceived), and background (that which is sensed, but not perceived). Thinking this way allows more in-depth consideration of this problem: questions such as 'what information is represented in the foreground (sensed), what is in the background (perceived)?'; 'can a transition between these states be achieved (and how)?'; 'is there too much visual data?'; or 'what information best suits which sense?' can be posed. Questions like this force an investigation into the quantity, and context of information that being communicated to the user, and also to consider the user's attention as a resource. These questions fit very well with the themes that reoccur throughout the previous project chapters, as is discussed later, and as such it seems that the consideration of senses, the information being conveyed, and how much attention is required from the user should be a key part of a design decision when working within this space.

The concepts of senses, Information Bandwidth, and Information Overload are further explored in the final project discussed within this thesis: the Antus (ANTUS) project. The previous project in the iterative design cycle, MF, highlighted some interesting concepts about the senses, the representation of information and its limits. In order to investigate these concepts further, ANTUS is a game designed specifically to present so much visual information to the players of the game that it would surpass that bandwidth limit of the visual channel and cause Information Overload. This game is used as a prop to introduce people to the ideas that have been discussed until this point, and a workshop was used to generate ideas for techniques to address information overload. Similar to the findings and insights discussed in the other projects, some of the findings were tightly coupled with the ANTUS project itself (such as the expectation for robot actors to have a personality is discussed more in Chapter 5), whilst others were more generalisable in nature. One of the most interesting findings is the way that people managed their attention when designing new information feedback systems: they often preferred "glanceable" feedback over those which required constant monitoring, employing attention like it is a resource, and using it sparingly. This is logical, and is supported by the previous works discussed, as well as other academic works in the area.

These themes and ideas can be used to construct the start of a framework for designers working within this space. This framework focuses on information, the senses, and the way information is conveyed and received by those interacting. The easiest way for these ideas to be presented are as an augmentation to Verplank's do?know?feel framework. As argued in Chapter 1, Verplank's do?Know?Feel? cycle outlined in his Interaction Design Sketchbook offers a flexible framework for understanding the way that people interact with the world around them. Although it was initially designed for those working with physical products, this flexibility allows this framework to be applicable to other points on the virtuality continuum. Considering that senses are the focus of this augmentation, this shall be referred to as the "Do? Know? Sense?" framework. This framework incorporates all of the elements of Verplank's framework as previously discussed, such as map and path, hot and cool media, but further defines the "Feel" portion of the framework by adding additional considerations for the designer. The framework considers how the user *feels*, or rather *senses* their environment. The concepts that of which this framework consists can be more succinctly represented as laid out in the following sections.

6.1.1 Senses

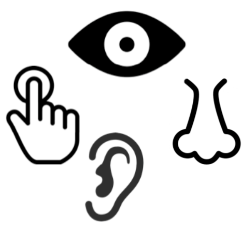

Figure 6.2: An icon representing the additional consideration of senses

A designer should consider how information is represented to the user. What sense will is be picked up on? Does this sense suit the information?

6.1.2 Bandwidth of senses & Information Overload

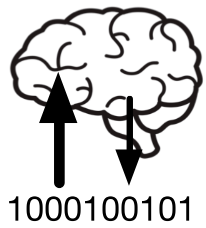

Figure 6.3: An icon representing the additional consideration of Information Bandwidth

How much information is being presented? Is there so much information that the realistic bandwidth of the sense is surpassed? Could this be alleviated representing it another way?

6.1.3 Attention as resource: Foreground and Background, Perceived or Sensed?

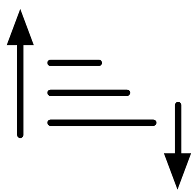

Figure 6.4: An icon representing the additional consideration of Attention as a Resource

Is this information in the foreground or background of the user's attention? Does it need to be where it is? Would it require less attention were it represented using another sense?

6.1.4 Bringing it all together

Bringing all of these concepts together with Verplank's original diagram, we can produce the cycle shown in Figure 6.5.

In this Do? Know? Sense? framework, the hybrid physical/virtual world is positioned at the centre, with the continuous cycle of Do, Know, Sense representing how these hybrid spaces are interacted with and understood. The elements that make up the "Know" phase of this cycle are unchanged from Verplank's original framing as maps & paths still seem valid with respect to the themes generated throughout this work. "Do" has been updated to represent both physical inputs (handles and buttons) but also virtual inputs: these have been represented by some common GUI widgets. The bulk of the observations throughout this work are represented to augmentations to the "sense" portion of this cycle. As it is an extension of Verplank's "Feel", it still considers different types of media (hot and cool). However, this framework also suggests that the way that the person receives information about the world is important. How did they sense this information? Was it consciously perceived or just sensed? How much, and in what form is this information?

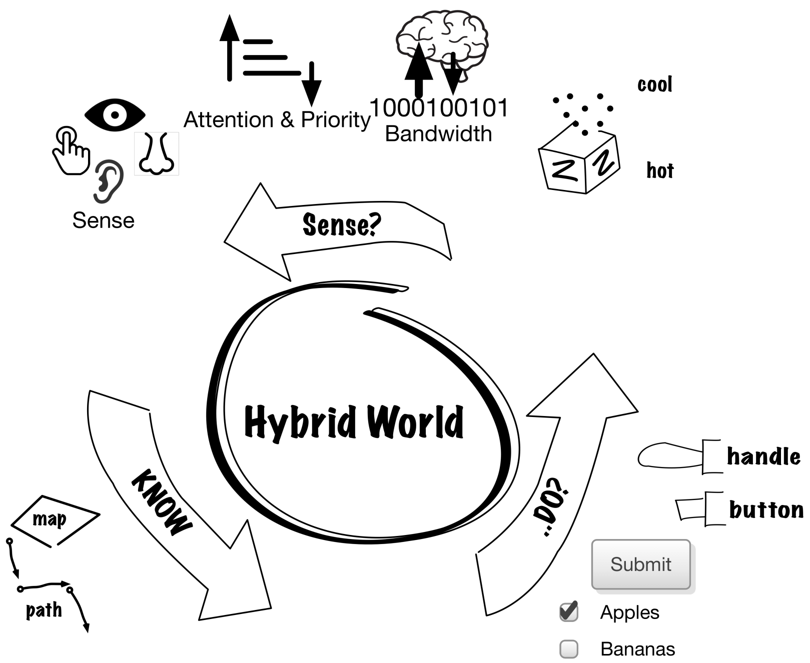

Figure 6.5: The Do? Know? Sense? diagram

Whilst this framework is by no means meant to be a complete and concrete framework that dictates the way that designers should design for interactions in this space, it is presented as a set of themes that designers can consider, and a framing for the flow of information in this space.

6.2 Meeting the Aims and Objectives

Chapter 1 discusses the aims and objectives as they were at the outset of this work. Therefore it is possible to reflect upon the aims and objectives and consider to what degree these were met. Chapter 1 outlines these aims and objectives as follows:

Aim:

To produce a framework that contributes to an improved ability to conceptualise Mixed Physical/Digital spaces.

Objectives:

- Investigate and explain the epistemological standing for creating a framework.
- Explore the problem space considering the aforementioned epistemological standing through Research *through* Design.
- Interpret the findings from the exploration into a useable framework.
- Provide reflections upon this process, and the produced framework.

Reflecting upon these aims and objectives suggests that this project has succeeded in reaching them. However, before considering these high level overarching aims and objectives the degree to which this work met the requirements for successful Research through Design should be considered. Chapter 2 outlines these, each of which is addressed here:

Documentation. At all stages of the process, I took care to ensure that sufficient documentation was captured for two reasons: it is needed at each stage for inductive analysis, and it is required both for portfolio format research, and for the methodology. I made it a priority to take pictures at most stages of the work so that it would be easier to later refer back. However, I found this task more difficult that previously imagined. Specifically, I found it difficult to remember to document failed avenues of research in detail. A good example is the IR tags discussed in the DTF project. I tried creating these tags before eventually settling on AR markers as a replacement, however, as is reflected in much of academia, it is often hard to document failure. I retrospectively took pictures of failed or discontinued avenues of research throughout the process of writing up but I do recognise that it would have been better to document these shortcomings at the time.

Connectivity. Discussion with the outside world is of great importance when differentiating 'little r' research from 'big R' research in the design space. I have always had this as a key goal of this work and as such, connectivity was achieved firstly through user studies of varying formality, and secondly via academic publication. Most of the ideas and prototypes in this thesis have been presented, published, and peer reviewed and the subsequent feedback has been influential on the direction of the work. The remaining works are currently in a state of works-in-progress to be published at a later date.

Analysis. Whilst Analysis is a very important part of research within the methodology outlined, this goal is almost inherent given that reflection is a core part of the process discussed. As such, I feel that analysis has been carried out in a reflective manner throughout the entirety of this thesis.

It seems clear that this work meets these targets well. Despite this, I think that a few components of the research process are not well captured by the previously presented process diagram. Particularly the way that each stage can produce multiple ideas and findings which may, or may not be relevant to the overall theories or research area. However, it would be difficult, and even unwise, for the researcher to decide which category the outcomes of each study fall in to for the same reasons it is difficult for the designer to be specific about the research question prior to beginning research. To better represent the actual process, I suggest amending the diagram as follows in Figure 6.6:

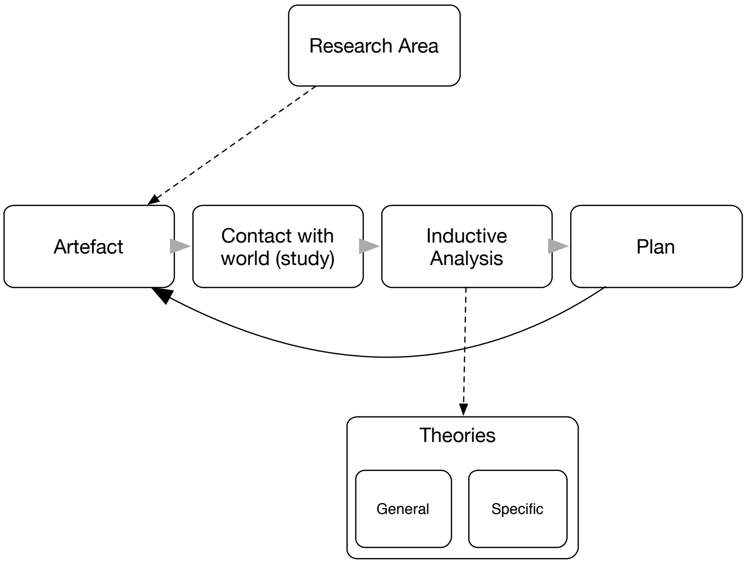

Figure 6.6: The modified research process diagram.

Specifically, the resulting process, complete with findings that contributed to the overall theories generated throughout this work can now be represented as per Figure 6.7.

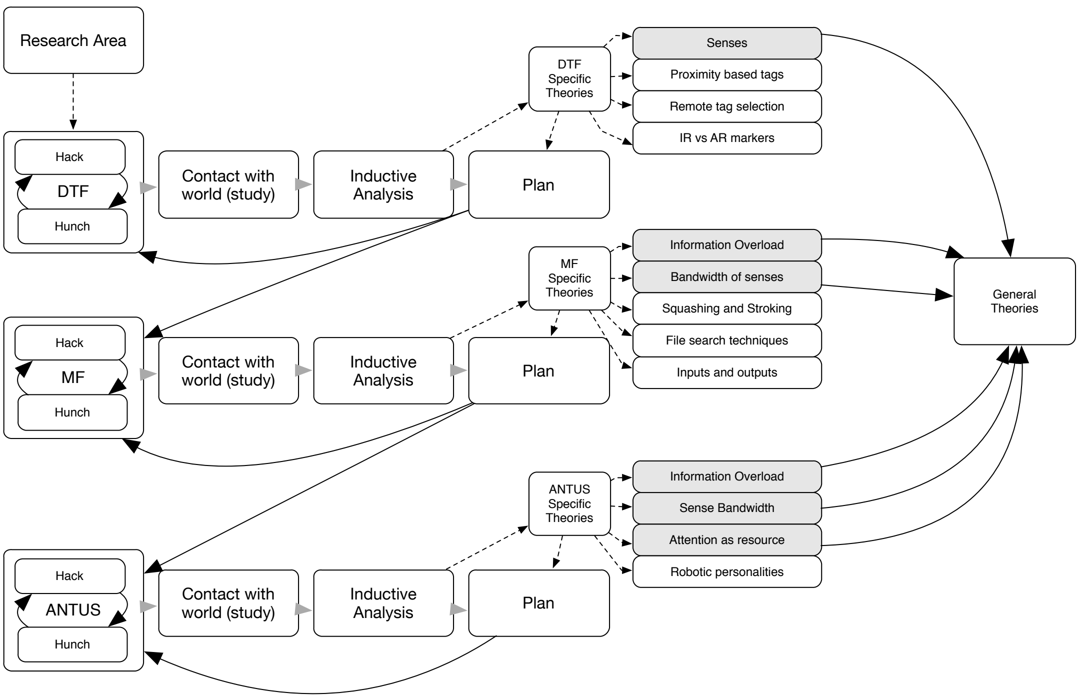

Figure 6.7: Specific process diagram.

This diagram shows well the way in which a research area can be narrowed to a more specific question through the act of research through design. Each stage of the iterative design process contributed to the theories that have been used to construct this framework, but at each stage, additional relevant information and theories are generated.

Following on from this it is possible to evaluate the degree in which this work met the higher level aims and objectives.

The first objective; "Investigate and explain the epistemological standing for creating a framework" is addressed in Chapter 2 which begins with an in-depth exploration and definition of epistemological positioning, relevant terminology (e.g Research through Design) as well as outlining contributing and parallel areas of work. At its conclusion, Chapter 2 outlines set of rules to follow when exploring the problem space.

The second Objective; "Explore the problem space considering the aforementioned epistemological standing through Research *through* Design" is addressed by Chapters 3, 4.1 and 5 which describe three iterative research projects:(DTF, Magnetic Files and ANTUS respectively). Throughout these projects interactions with the outside world highlighted themes for later research. Where possible the following iteration was planned and adapted in such a way that these themes could be explored.

The third objective; "Interpret the findings from the exploration into a useable framework" is tackled with the production of the "Do? Know? Sense?" framework discussed through much of the current chapter, but culminated in Section 6.1.4. This framework builds upon Verplank's "Do?Know?Feel" framework by augmenting it with concepts the more strongly relate to the hybrid physical/digital space in which this work is carried out.

The final objective; "Provide reflections upon this process, and the produced framework" is offered in Section 6.4.

Therefore, the overall Aim of the work - To produce a framework that contributes to an improved ability to conceptualise Mixed Physical/Digital spaces - can be considered a success.

6.3 Key contributions and outcomes of this work

Throughout this chapter I have discussed the outcomes of this work in a way that can be implicitly interpreted as the findings and outcomes of the research carried out throughout this project. However, I feel it is important to explicitly highlight these contributions too: some of these contributions, such as the methodology for this research, should be considered a contribution in itself. The following list contains the contributions of this work that I consider most important:

The Methodology described at the start of this work (see Chapter 2) outlines a research methodology focused on many iterative cycles of research that guide research over time. Whilst it seems clear from my experience that this is how much of work in the field of design is carried out, I found it particularly difficult to pin down an exact definition of the methodology, and epistemology to which a researcher may prescribe in order to understand the work that is done in design as explicit research. As a result, I highlight the most applicable epistemology for this form of research, and construct a hybrid methodology that best reflects this epistemology. Whilst this chapter is of importance to this work, I whole-heartedly believe that it is also of some benefit to other researchers in the area. As such I consider the whole chapter a key contribution of this work.

Academic Papers were a key goal throughout this research work. I believe that the production of academic papers is an important part of the research process for two reasons: firstly, they expose work to scrutiny from other academics, and secondly they act as a means of dissemination to the greater academic community. It is for these reasons that throughout the research and writing up of this thesis, I have had four academic papers directly or tangentially related to this thesis accepted by journals and/or conferences, with others awaiting review. The main works carried out have been submitted and presented at conferences: the most recent being the presentation of *Designing Information Feedback Within Physical/Digital Game Spaces* at DRS2016, which was well received. Whilst the middle chapters of this work are expanded versions of these papers, I consider the papers themselves to be key contributions to this work, as they will likely have the most impact on the wider academic community.

How do we do? How do we know? how do we *sense*? I would consider this, alongside the methodology, to be the most important contribution that I am trying to make with this work. Whilst the methodology chapter addresses the fact that no real explicit definition for methodology exists for work carried out in this style, I hope to also contribute some theory to build on existing relevant work that addresses the way in which we understand to interact with objects. More explicitly, the way that 'sensing' and 'perceiving' are an important part of the interaction process, and how current understanding for how we interact relies on the user knowing at what point along the virtuality continuum an object exists in order to interact with it. As such we have two streams of research: understanding how to interact with the physical, and understanding how to interact with the digital. As things are increasingly becoming both digital and physical, this categorisation is becoming less appropriate. I propose that the senses be used as a common area for the interaction process. An ideal way for this small idea to make the most sense is to replace 'How do we do? How do we know? how do we *feel*?' with 'How do we do? How do we know? how do we *sense*?'. This takes a framework that was designed for application in the physical reality and appropriates it to the hybrid space that is increasingly common.

In addition to this, a number of secondary contributions also come under this 'sense' banner: the concepts of 'channels', 'information overload', and 'attention as a resource' should also be considered contributions of this work.

Reflections , as highlighted by the methodology chapter are an intrinsic part of the methodology proposed: this work needs to be documented and reflected upon. This reflection comes in two forms: a reflection upon the work and findings within the projects, and a more overarching review of the methodology.

Documentation is very important to this research & methodology in a similar way to reflections. Whilst arguably a secondary contribution, I have included digital documentation of much of the work discussed throughout this thesis. This included digital copies of the papers submitted, source code for any software written, extra imagery taken throughout the studies discussed and tangential projects.

In summary, the key contributions of this work can be considered as the explicit methodology and epistemological positioning outlined in Chapter 2, a portfolio of projects (Chapters 3, 4.1 and 5) with respective academic papers, the "Do? Know? Sense?" framework outlined in Chapter 6 alongside relevant documentation and reflections upon the project.

6.4 A Personal Reflection on the Methodology

This thesis is presented as a portfolio, the arguments for which have made clear in Chapter 2. It follows that my reflections on this process are included and considered in part as a contribution. Similarly, the methodology employed in this work was constructed using key parts of other ways of knowing in order to fully respect my epistemological positioning as a researcher. Therefore it has not explicitly been used before and so my reflections upon its inaugural usage hold particular importance. Despite this, I believe that the methodology for this work is a somewhat formal framing for the existing design process. As such it can be argued that this has been very well tried and tested.

In the time before the commencement of this research work, the majority of research work I had been involved in had been carried out within the field of computer science. As such I was introduced to research from a somewhat positivist research perspective that focused primarily on the collection and dissemination of quantitative data - a positioning very different from the one displayed in this work. At the time I had no concept that alternative ways of doing research existed, "epistemology" was not a word in my vocabulary, and I would often be sceptical of work with data not represented in a table or a graph, or with a sample size too small to result in "statistically significant findings". It was therefore quite a research culture-shock when I started my training programme at Highwire, Lancaster university. Alongside other new researchers from varying fields (design, business, and even artists) I was asked to collaborate on tasks well outside of my expertise. It was immediately clear that everyone had their own process for gaining knowledge and even different criteria for what form this knowledge would take. If the task was to assess an interaction between people and an object I would consider the correct process to involve collecting data by asking multiple participants to perform the interaction task whilst some metric was recorded, and subsequently providing a yes/no or 1-5 style questionnaire after completion of the task. When comparing results with the other researchers the differences in the form of the collected data, and indeed what knowledge could be obtained from each became obvious. My results would offer findings such as "the proposed technique A task was completed 15% faster than technique B". At the time I felt confident that this meant technique A was an improvement over task B - my sample size of 30 participants was convincing. In contrast, other researchers would offer findings such as "the texture of the surface is unpleasant for the users", and "it feels unintuitive". My more positivistic approach, whilst producing interesting and valid data with respect to the speed of the two tasks, had limited the scope of the research to the questions I was asking "How fast is A compared to B?". Other approaches it seems have obtained answers to questions that I would have not thought to ask. It was at this point that I began to really think about what it was that I was I was measuring, and what I considered as Knowledge.

As discussed in Chapter 1, computer science, and specifically Human Computer Interaction (HCI), started out as a field that aimed to improve the usability of computers. In the early 1970s these areas were more similar to the field of maths and physics than the field it is today. As such, a positivist approach is logical: researchers knew exactly what they were trying to measure, and largely what the variables were. Since then, as already highlighted in Chapter 1, the field has grown much more complex and now incorporates many other fields in search of a greater understanding of the interaction process between human and computer. Personally, I think that the question of "usability" in this context now incorporates so many fields and angles that it is difficult to know where to start, especially when adopting a positivist approach. Throughout my time at Highwire, I gained more and more experience working in a designerly fashion. It became increasingly apparent to me that in order for a positivist approach to be successful in a field such as HCI, the researcher almost needed to be in a paradoxical state in which they have yet to do the study, but already knew the outcomes of the work in order determine the best line of questioning for research. The questions asked would heavily influence the direction of the research work. Designers, on the other hand had a less daunting task as it seemed the norm to their academic group to begin research with a direction, or research area rather than a clear, well defined research question. It was accepted that the designers did not yet know the answers and therefore were not in a position to dictate research questions. Instead, an iterative cycle of prototypes and discussions are used to continually re-evaluate, re-define, and concentrate the direction of research. This immediately appealed to me as a researcher as I had always struggled with planning research questions for this very reason.

To further my understanding of this process, I sought to find a formal description of the design process. However, although I found bits and pieces scattered about within the relevant literature, I could not find any formal description of the design process that described it in terms to which I could relate. The closest representations I could find were akin to Verplank's interaction design sketchbook which do highlight the process from a more applied perspective, centring on the iterative process and the way that a 'hack & hunch' cycle is often used. Yet I could not find anything that positioned this adequately with my epistemological preferences, and with reference to relevant methods for gaining knowledge. As a result, I pieced together a methodology that best reflects my experience of the design process, supported by relevant literature where possible, which is drawn together in Chapter 2.

My experience in using this methodology as the basis for this work is largely positive and consequently I consider myself a convert to the 'designerly' way of doing things. The process described in Chapter 2 discusses not only the theoretical underpinning that makes up this methodology, but also what the use of this methodology means in practice: starting with a research area, undergoing a number of iterative research cycles that encompass the artefact creation process, contact with the world, inductive analysis, plan, and then repeat. Finally, after a number of iterations, the researcher has gained enough knowledge about emergent themes to begin to inductively analyse the process, and produce theories based upon these themes. This process has tightly directed the course of work throughout this research. I started with the research area defined in Chapter 1. From here, I begin with the first iteration of the design cycle (the DTF project discussed in Chapter 3). This iteration incorporated the 'hack and hunch' process also discussed in Chapter 2 to build the projector/depth sensing system (artefact) used throughout the process. From here, the DTF system had contact with the outside world in the form of informal user studies. The discussions with participants then led to new idea generation via inductive analysis. In the case of DTF, these themes could be separated into those that were about usability of the DTF system specifically: such as dealing with 'bridges' with only one depth sensor, or working well with sand; and those that were more reflective of the research space as a whole: such as the idea that virtual elements could be sensed across multiple senses, rather than just being seen or heard. These themes were investigated and used to plan the next stage of the research process. This next stage was the MF project, which again started with artefact creation. In this case, the artefact borrowed a lot from the previous artefact (DTF) using a similar projector, Kinect, and AR markers to produce a space in which both virtual and physical elements exist. However, in this iteration, a focus was placed upon other ways to represent information, such as touch.

Working in this way meant that I had the flexibility to adapt and change the direction of the research as more knowledge was gained about the subject area. The knowledge gained was in a richer, albeit more varied form which lends itself to this type of more exploratory work.

Despite these positive reflections, there also seem to be a number of caveats to this research style. One of the only criticisms I have of this style is the lack of an explicit literature review stage. Although The reasoning behind this is argued and explained in Chapter 2, I feel that it perhaps puts the researcher at a disadvantage at the start of the process: without knowing what has been done in the area the research is at risk of "re-inventing the wheel". I would be interested to explore to what extent some level of literature review (that is carried out before the work rather than in this case carried out during) can be incorporated without detracting from the flexibility that this methodology offers.

Overall, I believe that this research style offers benefit to the researcher - particularly in those spaces that are rapidly evolving or immature. The acceptance and acknowledgement of the researcher (or designer) as a part of the process allows free exploration without placing too much emphasis on reducing "Mess". Similarly accepted is the concept that new knowledge gained throughout the process influences the direction of the research. This allows the researcher freedom to iterate and change the research throughout. I hope that alongside the framework this work outlines a methodology that others can adopt and build upon for their work.

6.5 Future Work

The work described and discussed throughout this thesis are inherently exploratory in nature, the reasons for which have been extensively covered. In addition to the work that has been discussed, there are a number of avenues that would be worth investigation given additional time and funding.

Up to this point, this work has exposed a number of themes that should be considered by designers when working in this hybrid space. I hope that one outcome of this work is that designers have a more formalised way to consider the exchange of information, and how people interact within this hybrid space. Having said this, each of these areas of consideration (such as Information Bandwidth, attention, and the use of senses) leads to a number of follow-on questions which could (and should) be answered.

One future project would be to investigate more explicitly the way in which people choose to use digitally driven sensory modules in every day life. This is actually a project that I started work on but later stopped in order to focus on writing up the projects in this thesis. For this I planned to use a number of "sense modules" which could communicate with a computer. Examples of these include a heat-mat, a smell dispenser, a servo-driven actuator, a speaker.

If a number of these were dispensed to people, and used in conjunction with a tool such as If This Then That (IFTTT)[1], people could make complex causal relationships: e.g "If I have a new email from Dan, Then make the smell modules smell like Roses", or "If the weather is going to be bad, then the heat module should be cold". Although in previous works I have looked at exploring the uses of these different scenarios, it would be interesting to see what people do with these modules: perhaps there are senses which are predominantly used for certain types of information, or people prioritise some senses over others. A project such as this would certainly go someway towards making the theories highlighted within this exploratory research more supported.

Similarly, more explicit studies on the senses and attention would be an obvious next step. Whilst I have discussed senses and the different attention requirements of them, there are obvious unanswered questions at this point. For example: What differentiates a low attention noise to a high attention noise? or smell? The smell of fire is surely more attention grabbing that some other smells for example. How can a transition between such attentions be achieved? For this, a more controlled, lab-based study may be more appropriate as the questions are more explicit.

There is also additional scope for work further investigating the limits of Information Bandwidths and their channels. A project of mine not included in the main chapters of this thesis involved the creation of a large (2m x 1.5m) LED matrix board to be used as an ultra low resolution screen (18 x 28 pixels) for the "Light Up Lancaster" festival which is a three day festival of lights held in the centre of Lancaster every year. On this board a space game was projected, which was controlled using two Leap Motion controllers. Obviously working with this low resolution limits the amount of information that can be displayed. Whilst not the topic of this work, this leads to the question of how little information can be displayed on this screen, and it still be considered a game. As technology has developed, game designers have seemingly always pushed for higher resolution screens and games. I think a valid line of questioning involves this screen resolution (and therefore the Information Bandwidth) as a conscious design choice rather than an assumption, and what the effects of these decisions entail.

[1] https://ifttt.com/

6.6 Conclusion

This work discusses the future of interaction spaces: from historically separate physical and digital spaces into hybrid spaces that express properties of both. As this space is somewhat immature in nature with no clear definitions for interaction within this space, I present a ground-up framework, complete with my epistemological standing, and accompanying methodology which closely reflects that of the design process for investigation into this space. Given this process, I investigate this space through a number of iterative Research through Design projects, and highlight some re-occurring themes. These themes are drawn together into the a framework allowing clearer understanding of the space, which I hope to lay the groundwork for future exploration of this area. This framework is largely focused around the senses, A logical focus given the disparity between how the senses are used in each space if considered separately.

In conclusion, the framework introduced by this thesis is one step towards a clear understanding of how interactions occur within these hybrid spaces. Given the trajectory of current technology it can be expected that these kinds of interactions are to become more common. As with the airport billboard scenario discussed in the introduction to this work, I don't think it necessary that designers always adhere religiously to the suggestions made by frameworks: there may be perfectly legitimate reasons for deciding otherwise. Having said this, I think it is important that the designers be given the choice to decide: a choice which is only possible by the existence of such frameworks. The framework that this work builds towards gives the designer the option to consider the senses and the flow of information in more detail as wanted by the designer, and I hope it goes some way towards helping work within this space.

List of Acronyms

TEI Tangible, Embedded, and embodied Interaction

DRS Design Research Society

HCI Human Computer Interaction

TUI Tangible User Interface

GUI Graphical User Interface

NUI Natural User Interface

AR Augmented Reality

VR Virtual Reality

DTF Digital Terra Firma

IR Infra Red

UBICOMP Ubiquitous Computing

WIMP Windows Icons Mouse Pointer

IoT Internet of Things

DTF Digital Terra Firma

URP Urban Planning and Design

QR Quick Response

MF MagneticFiles

ENIAC Electronic Integrator and Computer

LCD Liquid Crystal Display

IFTTT If This Then That

EPSRC Engineering and Physical Sciences Research Council

RGB Red, Green, Blue: a system for representing the colors to be used in an image

DIY Do It Yourself

CX Creative eXchange

API Application Programming Interface

CAD Computer-Aided Design

HFS Hierarchical File System

HFS+ Hierarchical File System Plus

XML eXtensible Markup Language

PDF Portable Document Format

USB Universal Serial Bus

LiPo Lithium Polymer

ANTUS Antus

LED Light Emitting Diode

IDE Integrated Development Environment

NPC Non-Player Character

OpenGL Open Graphics Library

3D Three Dimensional

Appendix A

Appx

A.1 Additional related work

The chapters discussed throughout this thesis cover the main outputs and discussion topic investigated. However, during this time I was also part of additional works that are related, yet did not exist as part of the iterative design process:

Game design in an Internet of things

(Coulton et al., 2014): A discussion on the role of game objects in the Internet of Things (IoT). This work investigating the changing role of digitally enabled objects within game spaces, and as such, played a role in the conceptualisation of the problem discussed through this thesis.

abstract: *Whilst no consensus yet exists on how the Internet of Things will be realised, a global infrastructure of networked physical objects that are readable, recognisable, locatable, addressable and controllable is undoubtedly a compelling vision. Although many implementations of the Internet of Things have presented these objects in a largely ambient sensing role, or providing some form of remote access/control, in this paper we consider the emerging convergence between games and the Internet of Things. This can be seen in a growing number of games that use objects as physical game pieces to enhance the players interaction with virtual games. These hybrid physical/digital objects present game* *designers with number of interesting challenges as they i) blur the boundaries between toys and games; ii) provide opportunities for free-form physical play outside the virtual game; and iii) create new requirements for interaction design, in that they utilise design techniques from both product design and computer interface design. Whilst in the past the manufacturing costs of such game objects would preclude their use within games from small independent games developers, the advent of low cost 3D printing and open software and hardware platforms, which are the enablers of the Internet of Things, means this is no longer the case. However, in order to maximise this opportunity game designers will need to develop new approaches to the design of their games and in this paper we highlight the design sensibilities required if they are to combine the digital and physical affordances within the design of such objects to produce good player experiences.*

Visual abstraction for games on large public displays

(Gullick et al., 2017) An investigation into the level of detail needed for a game on large public displays. This builds upon concepts discussed throughout this thesis.

Abstract: *From its earliest developments video game design has arguably been closely coupled to technological evolution particularly in relation to graphics. In very early games the limitations of technology led to highly abstracted graphics but as technology improved, abstraction has largely been left behind as developers strive towards ever-greater realism. Thus, games are generally drawing from conventions established in the mediums of film and television, and potentially limiting themselves from the possibilities abstraction may offer. In this research, we consider whether highly abstracted graphics are perceived as detrimental to gameplay and learnability by current gamers through the creation of a game using very low-resolution display that would accommodate a range of display options in a playable city. The results of trialing the game at a citywide light festival event where it was played by over 150 people indicated that abstraction made little difference to their sense of engagement with the game, however it did foster communication between players and suggests abstraction is a viable game design option for playable city displays.*

A.2 Code for projects

In addition to the included CD, I have uploaded the considerable amount of computer code used throughout these projects to a digital repository. Please visit these links for digital copies of the work, and instructions on installation and setup. For these to work you will typically need to use Eclipse and import the projects, Ubuntu Desktop (I used version 16.04).

DTF Project

https://gitlab.com/Gullick/DTF

MF Project

https://gitlab.com/Gullick/MF

ANTUS Project

https://gitlab.com/Gullick/ANTUS

References

A. Agarawala and R. Balakrishnan. Keepin'it real: pushing the desktop metaphor with physics, piles and the pen. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1283–1292. ACM, 2006. 65

K. Agnew. The spitfire: legend or history? an argument for a new research culture in design. *Journal of Design History*, 6(2):121–130, 1993. 25

K. Ashton. That internet of things thing. *RFiD Journal*, 22(7):97–114, 2009. 22

S. Bakker, D. Vorstenbosch, E. van den Hoven, G. Hollemans, and T. Bergman. Tangible interaction in tabletop games: studying iconic and symbolic play pieces. In *Proceedings of the international conference on Advances in computer entertainment technology*, pages 163–170. ACM, 2007. 103

J. Balbi. Epistemological and theoretical foundations of constructivist cognitive therapies: Post-rationalist developments. *Dialogues in Philosophy, Mental and Neuro Sciences*, 1(1):15–27, 2008. 26

E. I. Barakova and T. Lourens. Expressing and interpreting emotional movements in social games with robots. *Personal and ubiquitous computing*, 14(5):457–467, 2010. 117

S. Bloehdorn, O. Görlitz, S. Schenk, M. Völkel, et al. Tagfs-tag semantics for hierarchical file systems. In *Proceedings of the 6th International Conference on Knowledge Management (I-KNOW 06), Graz, Austria*, volume 8, 2006. 64

H. Bradbury-Huang et al. Transforming the generation and application of knowledge: A manifesto on quality in action research., 2009. URL `http://www.sagepub.com/journalsProdDesc.nav?prodId=Journal201642`. 28

G. Bradski. The opencv library. *Dr. Dobb's Journal of Software Tools*, 2000. 49

E. Brandt and C. Grunnet. Evoking the future: Drama and props in user centered design. In *Proceedings of Participatory Design Conference (PDC 2000)*, pages 11–20, 2000. 95

E. Brandt and J. Messeter. Facilitating collaboration through design games. In *Proceedings of the eighth conference on Participatory design: Artful integration: interweaving media, materials and practices-Volume 1*, pages 121–131. ACM, 2004. 95

M. Buchenau and J. F. Suri. Experience prototyping. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*, pages 424–433. ACM, 2000. 95

A. E. Clarke. Situational analyses: Grounded theory mapping after the postmodern turn. *Symbolic interaction*, 26(4):553–576, 2003. 29, 30

J. Corbin and A. Strauss. Basics of qualitative research 3rd edition, 2008. 29

J. M. Corbin and A. Strauss. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology*, 13(1):3–21, 1990. 29

P. Coulton. Sensing atoms and bits. *Sensory Arts and Design*, page 189, 2017. 16

P. Coulton, D. Burnett, A. Gradinar, D. Gullick, and E. Murphy. Game design in an internet of things. *Transactions of the Digital Games Research Association*, 1(3), 2014. 144

J. W. Creswell, V. L. Plano Clark, M. L. Gutmann, and W. E. Hanson. Advanced mixed methods research designs. *Handbook of mixed methods in social and behavioral research*, pages 209–240, 2003. 33

N. Cross. *Designerly ways of knowing.* Springer, 2006. 23

M. Crotty. *The foundations of social research: Meaning and perspective in the research process.* Sage, 1998. 26

N. Dahlbäck, A. Jönsson, and L. Ahrenberg. Wizard of oz studieswhy and how. *Knowledge-based systems*, 6(4):258–266, 1993. 95

N. K. Denzin and Y. S. Lincoln. *Handbook of qualitative research.* Sage publications, inc, 1994. 28

W. K. English, D. C. Engelbart, and M. L. Berman. Display-selection techniques for text manipulation. *IEEE Transactions on Human Factors in Electronics*, HFE-8, Issue: 1(1): 5–15, 1967. 22

A. Findeli. A quest for credibility: Doctoral education and research in design at the university of montreal. *Doctoral Education in Design*, pages 99–116, 1998. 24

G. W. Fitzmaurice, H. Ishii, and W. A. Buxton. Bricks: laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 442–449. ACM Press/Addison-Wesley Publishing Co., 1995. 9

S. Follmer, D. Leithinger, A. Olwal, A. Hogge, and H. Ishii. inform: dynamic physical affordances and constraints through shape and object actuation. In *Uist*, volume 13, pages 417–426, 2013. 11, 67

J. Forlizzi, C. DiSalvo, J. Bardzell, I. Koskinen, and S. Wensveen. Quality control: a panel on the critique and criticism of design research. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pages 823–826. ACM, 2011. 25

C. Frayling. Research in art and design. *London: Royal College of Art*, 1993. 23, 25

B. Gaver and J. Bowers. Annotated portfolios. *interactions*, 19(4):40–49, 2012. 25, 32

W. Gaver. What should we expect from research through design? In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 937–946. ACM, 2012. 25, 26

W. W. Gaver. Technology affordances. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 79–84. ACM, 1991. 6, 7

C. Ghaoui. *Encyclopedia of human computer interaction.* IGI Global, 2005. 22

J. J. Gibson. Perceiving, acting, and knowing: Toward an ecological psychology. *The Theory of Affordances*, pages 67–82, 1977. 6

B. G. Glaser and A. L. Strauss. *The discovery of grounded theory: Strategies for qualitative research.* Transaction publishers, 2009. 29

C. Gray and J. Malins. *Visualizing research: A guide to the research process in art and design.* Routledge, 2016. 33

B. M. Gross. *The managing of organizations: The administrative struggle*, volume 2. [New York]: Free Press of Glencoe, 1964. 96

D. Gullick and P. Coulton. Designing information feedback within hybrid physical/digital interactions. *Proceedings of Design Research Society Conference 2016*, 2016. 100

D. Gullick, P. Coulton, and M. Lau. Magnetic files: exploring tag based file systems using embodied files. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 613–617. ACM, 2015. 62

D. Gullick, D. Burnett, and P. Coulton. Visual abstraction for games on large public displays. In *Intelligent Technologies for Interactive Entertainment: 8th International Conference, INTETAIN 2016, Utrecht, The Netherlands, June 28–30, 2016, Revised Selected Papers*, pages 271–275. Springer, 2017. 145

J. Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 115–118. ACM, 2005. 49

J. Hardy and J. Alexander. Toolkit support for interactive projected displays. In *Proceedings of the 11th international conference on mobile and ubiquitous multimedia*, page 42. ACM, 2012. 42

J. Hare, S. Gill, G. Loudon, D. Ramduny-Ellis, and A. Dix. Physical fidelity: Exploring the importance of physicality on physical-digital conceptual prototyping. In *IFIP Conference on Human-Computer Interaction*, pages 217–230. Springer, 2009. 114

R. Harper, J. A. Hughes, and D. Z. Shapiro. Harmonious working and cscw: Computer technology and air traffic control. In *Studies in computer supported cooperative work*, pages 225–234. North-Holland Publishing Co., 1990. 8

E. Hoggan, S. A. Brewster, and J. Johnston. Investigating the effectiveness of tactile feedback for mobile touchscreens. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1573–1582. ACM, 2008. 4

V. D. Hopkin. Human factors implications of air traffic control automation. *ADVANCES IN HUMAN FACTORS ERGONOMICS*, 19:145–145, 1993. 8

E. Hornecker and J. Buur. Getting a grip on tangible interaction: a framework on physical space and social interaction. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 437–446. ACM, 2006. 19

H. Ishii and B. Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*, pages 234–241. ACM, 1997. 9, 10, 97

H. Ishii, C. Wisneski, S. Brave, A. Dahley, M. Gorbet, B. Ullmer, and P. Yarin. ambientroom: integrating ambient media with architectural space. In *CHI 98 conference summary on Human factors in computing systems*, pages 173–174. ACM, 1998. 10

H. Ishii, A. Mazalek, and J. Lee. Bottles as a minimal interface to access digital information. In *CHI'01 extended abstracts on Human factors in computing systems*, pages 187–188. ACM, 2001. 9

M. Jackson. Hand drumming. In L. Shawn, editor, *Things as they are: New directions in phenomenological anthropology*, chapter 9, pages 196–212. Georgetown University Press, 1996. 28

W. Jonas. Research through design through research-a problem statement and a conceptual sketch. In *Design Research Society International Conference, K. Friedman, T. Love, E. Côrte-Real, and C. Rust Eds., Lisbon*, 2006. 24

W. Jonas. Design research and its meaning to the methodological development of the discipline. *Design research now*, pages 187–206, 2007. 24

S. Jordà, G. Geiger, M. Alonso, and M. Kaltenbrunner. The reactable: exploring the synergy between live music performance and tabletop tangible interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 139–146. ACM, 2007. 11

J. Juul. *A casual revolution: Reinventing video games and their players.* MIT press, 2010. 15

A. Karypidis and S. Lalis. The tangible file system. In *Distributed Computing Systems Workshops, 2003. Proceedings. 23rd International Conference on*, pages 268–273. IEEE, 2003. 64

S. Kemmis, R. McTaggart, and R. Nixon. *The action research planner: Doing critical participatory action research.* Springer Science & Business Media, 2013. 28

M. Kojima, M. Sugimoto, A. Nakamura, M. Tomita, H. Nii, and M. Inami. Augmented coliseum: An augmented game environment with small vehicles. In *Horizontal Interactive Human-Computer Systems, 2006. TableTop 2006. First IEEE International Workshop on*, pages 6–pp. IEEE, 2006. 102

B. Koleva, S. Benford, K. H. Ng, and T. Rodden. A framework for tangible user interfaces. In *Physical Interaction (PI03) Workshop on Real World User Interfaces*, pages 46–50, 2003. 19

T. Költringer and T. Grechenig. Comparing the immediate usability of graffiti 2 and virtual keyboard. In *CHI'04 Extended Abstracts on Human Factors in Computing Systems*, pages 1175–1178. ACM, 2004. 11

G. Kortuem, F. Kawsar, V. Sundramoorthy, and D. Fitton. Smart objects as building blocks for the internet of things. *IEEE Internet Computing*, 14(1):44–51, 2010. 40

J. Law. *Making a mess with method*. Sage, 2007. 21, 31

A. V. Le, S.-W. Jung, and C. S. Won. Directional joint bilateral filter for depth images. *Sensors*, 14(7):11362–11378, 2014. 46

J. Leitner, M. Haller, K. Yun, W. Woo, M. Sugimoto, M. Inami, A. D. Cheok, and H. D. Been-Lirn. Physical interfaces for tabletop games. *Computers in Entertainment (CIE)*, 7(4): 61, 2009. 102

C. Lewis. *Using the" thinking-aloud" method in cognitive interface design*. IBM TJ Watson Research Center, 1982. 53

K. Lynch. *The image of the city*, volume 11. MIT press, 1960. 14, 124

W. E. MacKay. Is paper safer? the role of paper flight strips in air traffic control. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 6(4):311–340, 1999. 7, 8, 111

M. Maclure. Postmodernism: a postscript. *Theory and Practice in Action Research: some international perspectives*, page 61, 2002. 28

C. Magerkurth, M. Memisoglu, T. Engelke, and N. Streitz. Towards the next generation of tabletop gaming experiences. In *Proceedings of Graphics interface 2004*, pages 73–80. Canadian Human-Computer Communications Society, 2004. 102, 103

J. Marco, E. Cerezo, and S. Baldassarri. Toyvision: a toolkit for prototyping tabletop tangible games. In *Proceedings of the 4th ACM SIGCHI symposium on Engineering interactive computing systems*, pages 71–80. ACM, 2012. 114

J. McGrenere and W. Ho. Affordances: Clarifying and evolving a concept. In *Graphics interface*, volume 2000, pages 179–186, 2000. 7

M. McLuhan. *Understanding media: The extensions of man.* MIT press, 1994. 13, 124

R. McTaggart. 16 tenets of participatory action research. *Retrieved March*, 8:2005, 1989. 28

R. McTaggart. Principles for participatory action research. *Adult Education Quarterly*, 41(3): 168–187, 1991. 28

P. Milgram, H. Takemura, A. Utsumi, and F. Kishino. Augmented reality: A class of displays on the reality-virtuality continuum. In *Photonics for industrial applications*, pages 282–292. International Society for Optics and Photonics, 1995. 16

G. E. Moore et al. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85, 1998. 2

R. Niiyama, L. Yao, and H. Ishii. Weight and volume changing device with liquid metal transfer. In *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction*, pages 49–52. ACM, 2014. 81

D. A. Norman. Affordance, conventions, and design. *interactions*, 6(3):38–43, 1999. 6

D. A. Norman. The psychopathology of everyday things. *Foundations of cognitive psychology: core readings. MIT Press, Cambridge, MA*, pages 417–443, 2002. 6

M. Oakeshott. *Experience and its Modes.* Cambridge University Press, 2015. 28

OSXTags. OS X: Tags help you organize your files, 2015. URL `https://support.apple.com/en-gb/HT202754`. 65

ParrotPot. #CES 2015 : Parrot Pot, the most advanced connected plant pot, Jan. 2015. URL `http://blog.parrot.com/2015/01/05/ces-2015-flower-power-pot-the-smart-pot-that-grows-healthy-plant/`. 18

J. Patten and H. Ishii. Mechanical constraints as computational constraints in tabletop tangible interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 809–818. ACM, 2007. 11

J. Patten, H. Ishii, J. Hines, and G. Pangaro. Sensetable: a wireless object tracking platform for tangible user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 253–260. ACM, 2001. 49

J. Piaget. *Introduction à l'épistémologie génétique.* Presses Universitaires de France, 1950. 26

J. Piaget and E. Duckworth. Genetic epistemology. *American Behavioral Scientist*, 13(3): 459–480, 1970. 26

J. Piaget and P. Wells. *Psychology and epistemology: Towards a theory of knowledge.* Penguin Harmondsworth, 1972. 26

F. Preux. Rôle des strips dans l'activité des contrôleurs. *Sélection Professionnelle*, 1994. 8

N. Ralph, M. Birks, and Y. Chapman. The methodological dynamism of grounded theory. *International Journal of Qualitative Methods*, 14(4):1609406915611576, 2015. 29

R. Ramirez. An epistemology for research through design. In *Proceedings of the ICSID Design Education Conference*, pages 1–14, 2009. 27, 29, 30, 31

J. Reimer. A history of the gui. *Ars Technica*, 5, 2005. 22

H. W. Rittel and M. M. Webber. Dilemmas in a general theory of planning. *Policy sciences*, 4 (2):155–169, 1973. 22

P. Rogers, R. Puryear, and J. Root. Infobesity: The enemy of good decisions. *Bain and Company*, 2013. 96

D. A. Schön. *Educating the reflective practitioner: Toward a new design for teaching and learning in the professions.* Jossey-Bass, 1987. 28, 121

J. Schöning, J. Hook, T. Bartindale, D. Schmidt, P. Oliver, F. Echtler, N. Motamedi, P. Brandl, and U. von Zadow. Building interactive multi-touch surfaces. In *Tabletops-Horizontal Interactive Displays*, pages 27–49. Springer, 2010. 11

T. A. Schwandt. Constructivist, interpretivist approaches to human inquiry. In N. K. Denzin and Y. S. Lincoln, editors, *Handbook of qualitative research.*, pages 18–137. Sage Publications, Inc, 1994. 27

O. Shaer and E. Hornecker. Tangible user interfaces: past, present, and future directions. *Foundations and Trends in Human-Computer Interaction*, 3(1–2):1–137, 2010. 19

O. Shaer, N. Leland, E. H. Calvillo-Gamez, and R. J. Jacob. The tac paradigm: specifying tangible user interfaces. *Personal and Ubiquitous Computing*, 8(5):359–369, 2004. 19

C. Speier, J. S. Valacich, and I. Vessey. The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences*, 30(2):337–360, 1999. 96

S. L. Star and J. R. Griesemer. Institutional ecology,translations' and boundary objects: Amateurs and professionals in berkeley's museum of vertebrate zoology, 1907-39. *Social studies of science*, 19(3):387–420, 1989. 54

E. Stolterman. The nature of design practice and implications for interaction design research. *International Journal of Design*, 2(1), 2008. 25

A. Strauss, J. Corbin, et al. *Basics of qualitative research*, volume 15. Newbury Park, CA: Sage, 1990. 29

C. Thacker, E. McCreight, B. Lampson, R. Sproull, and D. Boggs. Alto: A personal computer, August 1979. URL `https://www.microsoft.com/en-us/research/publication/alto-personal-computer/`. 22

A. Toffler. *Future shock*. Bantam, 1971. 96

B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM systems journal*, 39(3.4):915–931, 2000. 12, 19

J. Underkoffler and H. Ishii. Urp: a luminous-tangible workbench for urban planning and design. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 386–393. ACM, 1999. 11, 48, 49

Unknown. Artoolkit, 2005a. URL `https://artoolkit.org/`. 50

Unknown. Nyartoolkit project, 2005b. URL `http://nyatla.jp/nyartoolkit/wp/?page_id=198`. 50

Unknown. Nyid marker maker, 2005c. URL `http://sixwish.jp/AR/Marker/idMarker/`. 50

Unknown. Artoolkit marker generator online released, 2014. URL `http://flash.tarotaro.org/blog/2008/12/14/artoolkit-marker-generator-online-released/`. 50

W. Verplank. Interaction design sketchbook by frameworks for designing interactive products and systems. *Last accessed*, 18(2):2014, 2009. 11, 12, 34

M. Waldner, J. Hauber, J. Zauner, M. Haller, and M. Billinghurst. Tangible tiles: design and evaluation of a tangible user interface in a collaborative tabletop setup. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments*, pages 151–158. ACM, 2006. 11

M. Weber. *On the methodology of the social sciences*. Lulu Press, Inc, 2015. 21

M. Weiser. Ubiquitous computing. In *ACM Conference on Computer Science*, page 418, 1994. 22

D. Wigdor and D. Wixon. *Brave NUI world: designing natural user interfaces for touch and gesture*. Elsevier, 2011. 10

A. Wilson, H. Benko, S. Izadi, and O. Hilliges. Steerable augmented reality with the beamatron. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, pages 413–422. ACM, 2012. 42

Z. Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2):4–10, 2012. 41

J. Zimmerman and J. Forlizzi. The role of design artifacts in design theory construction. *Artifact*, 2(1):41–45, 2008. 25

J. Zimmerman, J. Forlizzi, and S. Evenson. Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 493–502. ACM, 2007. 24, 25

J. Zimmerman, E. Stolterman, and J. Forlizzi. An analysis and critique of research through design: towards a formalization of a research approach. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, pages 310–319. ACM, 2010. 25