Technical Report

UCAM-CL-TR-733
ISSN 1476-2986

Number 733

UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Tangible user interfaces for peripheral interaction

Darren Edge

December 2008

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

This technical report is based on a dissertation submitted January 2008 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Jesus College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Tangible User Interfaces for Peripheral Interaction

Episodic Engagement with Objects of Physical, Digital & Social Significance

Darren Edge

Since Mark Weiser's vision of ubiquitous computing in 1988, many research efforts have been made to move computation away from the workstation and into the world. One such research area focuses on "Tangible" User Interfaces or TUIs – those that provide both physical representation and control of underlying digital information.

This dissertation describes how TUIs can support a "peripheral" style of interaction, in which users engage in short, dispersed episodes of low-attention interaction with digitally-augmented physical tokens. The application domain in which I develop this concept is the office context, where physical tokens can represent items of common interest to members of a team whose work is mutually interrelated, but predominantly performed independently by individuals at their desks.

An "analytic design process" is introduced as a way of developing TUI designs appropriate for their intended contexts of use. This process is then used to present the design of a bimanual desktop TUI that complements the existing workstation, and encourages peripheral interaction in parallel with workstation-intensive tasks. Implementation of a prototype TUI is then described, comprising "task" tokens for work-time management, "document" tokens for face-to-face sharing of collaborative documents, and "contact" tokens for awareness of other team members' status and workload. Finally, evaluation of this TUI is presented via description of its extended deployment in a real office context.

The main empirically-grounded results of this work are a categorisation of the different ways in which users can interact with physical tokens, and an identification of the qualities of peripheral interaction that differentiate it from other interaction styles. The foremost benefits of peripheral interaction were found to arise from the freedom with which tokens can be appropriated to create meaningful information structures of both cognitive and social significance, in the physical desktop environment and beyond.

Acknowledgements

I am deeply indebted to many people for the support they have given me throughout the course of this PhD. It would not have been possible without the financial support of The Boeing Company, whose investment in emerging technologies research at the University of Cambridge enabled me to pursue a rich and rewarding PhD in an area of great promise. Special thanks to Anthony Majoros of Boeing for overseeing the project.

My decision to undertake a PhD was in no small part due to my supervisor, Alan Blackwell, whose approach to research is both inspirational and humbling. His patient guidance and thoughtful advice have been instrumental in my development as a researcher, and I hope that I have justified the confidence he has shown in me since the beginning.

The photographs in this dissertation were expertly staged and shot by Luke Church, who also identified numerous areas for improvement in draft versions of this document. Chris Nash provided additional valuable feedback on dissertation drafts, as well as taking the time to pose for photographs depicting various aspects of interface usage. Phil Tuddenham and Luis Vargas were also kind enough to appear in these photographs, as well as devoting many hours to the discussion of my research throughout my PhD. As much as I appreciate these contributions of Luke, Chris, Phil and Luis, I value their friendship more so. Similarly for Lorisa Dubuc, Daniel Bernhardt, Will Billingsley, and Cecily Morrison – thank you all for making the Rainbow Group a bright and colourful place to be.

I am lucky enough to have had two families for as long as I can remember, and I would like to express my gratitude to all four of my parents for their enduring love and encouragement. Most of all, however, I am grateful to Emily. She has supported me for the entirety of my University life, and my life would not be complete without her. Thank you for everything.

For Emily

Contents

List of Tables

List of Figures

Chapter 1

Introduction

As the tasks for which we use our computers become ever more distributed and time-pressured, the need increases to find ways of interacting with digital information that are more expressive and direct than pointing-and-clicking. The conventional WIMP interaction paradigm has three primary objects of interest (*windows*, *icons* and *menus*), and a single point of control (the *pointer*) operated by a single device (commonly a mouse). Even when augmented with a keyboard, the expressive capabilities of this workstation arrangement are limited to just six fundamental actions: select, position, orient, path, quantify, and text (Foley et al., 1984). In comparison, our everyday interactions with the physical world are many and varied, leading to a level of perceptual awareness and motor control far in excess of anything required by conventional graphical interfaces. Utilising our spare capacity for interaction – through our dexterity, our bimanual skill, and our abilities of peripheral vision and spatial memory, amongst others – requires a reinterpretation of Shneiderman's (1987) theory of direct manipulation that is more substantial than the relatively *indirect* "direct manipulation" of WIMP interfaces. One such reinterpretation has given rise to the field of tangible interaction, in which Tangible User Interfaces, or TUIs, branch out into the real world by giving physical form to digital information and its subsequent control (Ishii and Ullmer, 1997).

Yet despite the apparent inadequacies of single-user workstations, they proliferate in almost all office-based work environments. The development of a multipurpose, workstation-replacement TUI is an unrealistic goal, since the necessary concreteness of physical representations cannot hope to deal with the infinite malleability of digital information. Similarly, the addition of extra display screens, keyboard hotkeys, dedicated pan-and-zoom devices and so on may improve the 'bandwidth' of interaction, but ultimately this is only a partial solution, still limited by the constraints of the WIMP paradigm. To be considered "tangible" in the sense accepted within the TUI community, the physical elements of the interface must be meaningful in themselves – as objects of representation, as well as control.

This defines *what* TUIs are, but it does not justify *why* they should be used in preference to other classes of interface, such as those based on speech or gesture recognition. One argument is that TUIs can take better advantage of a broader range of human abilities, but this line of reasoning is fraught with the burden of proof: which abilities are being used; is there a statistical and substantial improvement relative to other interfaces; and how generalisable is the effect? My preferred justification, therefore, is to say that *TUIs enable users to exploit and augment their spatial and social contexts, through the meaningful coupling of physical and digital media*. This incorporates the *what* and the *why*, and encourages a broader appreciation of both.

The original conception of tangibles, as "tangible bits" formed through fusion of the physical and the digital (Ishii and Ullmer, 1997), translated Weiser's vision of "calm" computing (Weiser and Brown, 1995) into "graspable media" on the one hand – located in the foreground of activity and at the focus of users' attention – and "ambient media" on the other, existing in the background of activity and at the periphery of users' attention. Whereas calm technology "engages both the center and the periphery of our attention, and in fact moves back and forth between the two" (ibid.), the historic distinction between graspable and ambient media reifies the notions of centre and periphery as fixed categories of the world, rather than treating them as transient states of the mind. In comparison, the style of TUI I present in this dissertation is a more faithful interpretation of calm computing, drifting between the focus and periphery of users' attention according to the demands of work activity. Interaction with the TUI is not session-based but *episodic* – composed of short "episodes" of use that are loosely related and dispersed in time. I use the term "peripheral interaction" to describe this new class of TUI: *'peripheral interaction' is about episodic engagement with tangibles, in which users perform fast, frequent interactions with physical objects on the periphery of their workspace, to create, inspect and update digital information which otherwise resides on the periphery of their attention.* This style of interaction achieves a balance between the "calm technology" of Weiser and the "engaging user experiences" of Rogers (2006), in which "people rather than computers [...] take the initiative to be constructive, creative and, ultimately, in control of their interactions with the world – in novel and extensive ways". When the user is immersed in their primary tasks, interaction with the objects of the interface can be intermittent and cursory, without paying much attention to their physical form, or indeed to the digital updates being made. When the user is detached from the demands of their tasks, they are free to engage more fully with the objects of the interface, taking time to customise or annotate their physical form, or using them to make more thoughtful updates to the underlying digital information.

In its most general sense as a category of interaction styles, "peripheral interaction" can be seen as any kind of interaction with objects – physical or digital – that do not occupy the typical centre of the user's attention. As such, occasional touch-based interactions with a secondary display surface could be considered "peripheral" relative to periods of focused concentration on the primary workstation. However, the particular quality of tangibles that make them appropriate for peripheral interaction is the fact that they create an interaction periphery beyond any fixed display surface or surfaces, extending the physical-digital frontier away from the user's normal centre of attention and into their spatial and social contexts. My definition of peripherality is therefore about breaking away from users' attentional and spatial norms of interaction. Tangibles can be seen as the purest objects of peripheral interaction, since they provide more 'things' to divert the user's attention away from its established focus, and more places in which to put them.

The target domain for the style of TUI described in this dissertation is the office context, in which the work of individuals is performed independently but as part of an greater team effort, typically orchestrated through weekly meetings. In a peripheral-interaction TUI for this context, items of interest to individuals or team members could be externalised as digitally-augmented physical-tokens, exploiting the positive characteristics of both physical and digital media. Such tokens could provide interactional benefits in terms of allowing both "at a glance" viewing of digital information and Heideggerian "ready to hand" availability of physical handles to that information. As such, *peripheral TUIs could reduce barriers to interaction by facilitating the performance of periodic, low-attention activities in parallel with workstation intensive tasks, preventing such 'auxiliary' work activities from becoming neglected or forgotten.* The interactional benefits of such an arrangement are predicated on the TUI supporting a complementary activity type to those requiring the existing workstation of monitor, mouse and keyboard.

I will argue during the course of this dissertation, however, that it is a mistake to fixate on the interactional potential of TUIs. Whilst TUIs could indeed use the "bandwidth of the fingers" as a way to "beat the mouse" for certain tasks (Card et al., 1991), their potential benefits extend beyond the purely interactional: TUIs should be viewed holistically as vehicles for the joint delivery of interactional, cognitive and social benefits. Individuals' desktop environments are often rich in meaning, having a "natural history" (Malone, 1983), and providing a habitat for their owners in which they perform situated actions (Suchman, 1987) based on evolving work situations. The configuration of physical and digital artefacts can also be seen as traces of previous external cognition (Preece et al., 2002), providing entry points for future work activities (Kirsh, 2001). In this context, *peripheral TUIs could take advantage of existing strategies for the projection of meaning onto the environment, by allowing users to structure their understanding of the current work situation via the meaningful arrangement of salient physical tokens in their physical desktop.*

Individuals rarely work autonomously and in isolation, so in addition to the cognitive impact of TUIs there are also likely to be social implications. The use of artefacts for the coordination of work efforts is addressed theoretically by the shift from external to distributed cognition (Hutchins, 1994), whilst the sociological theory of symbolic interactionism holds that a person's actions towards a thing are based on the symbolic meaning it has for them, derived from a history of social interaction with others with respect to the thing (Blumer, 1969). In the office context, the collective social benefits of peripheral interaction with physical tokens grow organically from their use as instruments of interaction and cognition. Firstly, the lowering of barriers through physical interaction has the side-effect of enabling digital updates to be published and communicated to interested parties in real-time. Secondly, the cognitive organisation of tokens in the environment has the side-effect of presenting a highly visible, symbolic embodiment of social roles and relationships to other group members and passers-by. Thirdly, the use of customisable physical tokens for interaction and cognition has the side-effect of providing a common unit of physical, digital and social exchange, encouraging face-to-face conversation when the TUI design requires that tokens be passed from one person to another. In summary, *peripheral-interaction TUIs could support new systems of symbolic social meaning within groups: created through the association of digital information with physical tokens; shared with others through the physical presentation of tokens and the digital publication of information; and modified through social interactions that are themselves mediated by token reference and exchange.*

These interactional, cognitive and social benefits of peripheral interaction are realised in my design of a bimanual desktop TUI, which incorporates an interactive surface and control device to support the peripheral inspection and manipulation of the digital information attributes associated with augmented physical tokens. This design complements the existing monitor, mouse and keyboard, and supports the auxiliary work activities of task management, document sharing, and status communication. The contextual motivation and evaluation of the design will be discussed in depth in later chapters. An overview of the thesis underlying this research is presented next.

1.1 Thesis Summary

In *peripheral interaction*, users are free to arrange independently-meaningful, digitally-augmented physical tokens on the periphery of their workspace and away from their normal centre of attention, ready to selectively and fluidly engage those tokens in loosely related, dispersed episodes of digital, cognitive, and social use. It has the potential to physically lower interaction barriers, digitally share the results of interactions, and symbolically engender new forms of social value. The key value of peripheral interaction is realised not so much in moments of interaction, but in the momentary shifts of attention afforded by having multiple tokens that are meaningful on multiple levels.

A *bimanual desktop TUI* supports the peripheral inspection and manipulation of digital information attributes associated with physical tokens. This interface structure is particularly suitable for the office context, complementing the existing monitor, mouse and keyboard, and supporting the performance of auxiliary work activities in parallel with primary workstation-intensive tasks. Such auxiliary activities are not just about the interactional demands of updating digital information, but about the cognitive and social demands of remembering what things to do, when to do them, and how to inform others about plans and progress. The cognitive and social implications of the TUI are that users can project their own systems of meaning onto both the desktop environment and the social exchanges of and around tokens, which represent the things of personal and collective value in the current work context.

1.2 Aims & Objectives

The aims and objectives of this work, determined at the start of the research project, were as follows:

1. To investigate design opportunities for the enhanced interaction modalities afforded by tangibility.
2. To augment the decision making capability of users who are distracted or impaired.

3. To design and implement a TUI prototype motivated by the needs of users in context.
4. To evaluate the TUI prototype during its situated use in an extended field deployment.
5. To contribute to the body of knowledge on TUI application, TUI design, and TUI evaluation.

1.3 Research Contributions

The major research contributions of this work to the field of HCI are as follows:

1. *TUI Application*. Identification of opportunities for low-attention interaction in office contexts, leading to a new class of desktop TUI that supports bimanual peripheral interaction.
2. *TUI Design*. A new process for the analytic design of contextually appropriate TUIs, supplemented by a specific design vocabulary for discussing physical token roles in TUI design.
3. *TUI Evaluation*. A new approach to the evaluation of TUI benefits in real contexts, leading to empirically grounded justifications for the choice of peripheral interaction.

1.4 Structure of Dissertation

This introductory chapter – Chapter 1 – has given an overview of the dissertation in terms of its theme, thesis, aims, and contributions. The chapter following on from this – Chapter 2 – presents a review of the TUI literature on which this dissertation builds, describing the conceptual evolution of the TUI research area, techniques used in TUI design, and published accounts of TUI implementation and evaluation. It concludes by summarising the general research themes within the literature, and identifying a research opportunity for personal, office-based TUIs.

Chapter 3 describes the analytic design of TUIs, bringing together perspectives from HCI and beyond, and formulating them into a coherent design process for the creation of TUIs that are appropriate to their usage context. In Chapter 4, this method is applied to the office context identified in Chapter 2 as lacking TUI support, and used to construct an argument for a "peripheral" style of interaction appropriate for the auxiliary activities of office life. The outcome of Chapter 4 is the abstract design of a peripheral interaction TUI, whose concrete design and implementation are presented in Chapter 5.

Chapters 6 and 7 present the field deployment of the prototype and its subsequent evaluation. Chapter 6 describes the way in which the fieldwork was set up in terms of methodology and methods, and uses log data and interview transcripts to form evaluations of the

TUI's impact on both individual and social work-spaces and work-practices. Having shown in Chapter 6 that the TUI supports peripheral interaction in the form of fast, frequent digital interactions, in Chapter 7 the broader benefits of peripheral interaction, as judged by the fieldwork participants, are examined more specifically using a card-sorting technique developed for this purpose.

The final, conclusion chapter – Chapter 8 – discusses the significance of this work in the context of HCI, reflects on peripheral interaction and the research questions posed in Chapter 4, proposes implications for the design of future TUIs, and sets an agenda for future work on both peripheral interaction and TUIs in general.

Chapter 2

Review of Tangible User Interfaces

The previous chapter introduced the general aims of this dissertation and its focus on the area of Tangible User Interfaces, or TUIs. In this chapter, I present a survey of the TUI literature, beginning with a discussion of the conceptual origins and evolution of the 'TUI' concept. I then go on to describe 'practical' TUI research: design methods for concept generation and analysis, a wide variety of TUI systems, and the various ways in which these TUIs have been evaluated. The chapter concludes with summary of research opportunities arising from analysis of the presented literature.

2.1 Conceptual Origins and Evolution of TUI Research

The idea of having human-computer interfaces with dedicated physical forms is not as recent as some research might suggest. It is only the revolution of personal desktop computing – with its inception in the Xerox Star (1976), and subsequent evolution through various iterations of PC and Apple based computers – that has allowed us to forget the custom-built hardware of the computing machinery that preceded the now pervasive monitor, mouse and keyboard. Many of the computers from the 1950s exposed much of their functionality through switches, knobs and dials, whilst the book *How to Build a Working Digital Computer* (Alcosser et al., 1967) described how a simple computer could be built from common household items, such as paperclips for switches and tin cans for drum memory. The modern PC represented a paradigm shift relative to these earlier "computing machines", and it would take another fundamental shift – courtesy of Mark Weiser's vision of ubiquitous computing in 1988 – to re-examine the concomitant roles of physicality and spatiality within the field of Human-Computer Interaction.

2.1.1 TUI Precursors: Ubiquitous Computing

In his seminal Scientific American article on ubiquitous computing – *The Computer for the 21st Century* – Weiser (1991) outlined "a new way of thinking about computers in the world, one that takes into account the natural human environment and allows the computers themselves to vanish into the background". This "disappearance" due to familiarity has been identified as something that allows users to think in terms of problems rather than tools; this phenomenon is referred to by prominent members of their disciplines (ibid.) as "compiling" (computer scientist, economist and Nobelist Herb Simon), "the tacit dimension" (philosopher Michael Polanyi), "visual invariants" (psychologist TK Gibson), "the horizon" (philosopher Georg Gadamer), "ready-to-hand" (philosopher Martin Heidegger), and "the periphery" (Xerox PARC computer scientist John Seely Brown). Whilst ubiquitous computing is about the perception of computer-mediated interaction fading into the background, it is not about "ambient" background computing; rather, it is about a blurring of the divide between the physical and the digital, such that we no longer notice the difference. Tangible computing is therefore just one realisation of the ubiquitous computing vision, making computation *more real* by augmenting the physical objects in the environment with digital information, and vice versa, until familiarity with such augmentations render them as the new reality in our subconsciousness.

In another of Mark Weiser's articles, coauthored with John Seely Brown (Weiser and Brown, 1995), *Designing Calm Technologies* highlights the distinction between designs that encalm and designs that inform. At the very heart of Information Technology, as conventionally received, is constant and continual competition for our one focus of attention; 'calm' technology, on the other hand, "engages both the *center* and the *periphery* of our attention, and in fact moves back and forth between the two" (ibid.). The periphery refers to aspects of our environment to which we are attuned, without demanding the focus of our attention. As such, we may not even be aware of our periphery until an unexpected event triggers a conscious response, such as becoming aware of a ticking clock only when it stops. Peripheral technologies are calming because the user regains control of information and the process of informing, selecting amongst many peripheral technologies for the next momentary focus. Information is pulled by the user from the environment as and when it is appropriate for them, rather than being burdened by information being pushed onto them by a range of technologies all competing for their attention. The principle that "we must learn to design for the periphery so that we can most fully command technology without being dominated by it" (ibid.) is central to my thesis that TUIs are a suitable vehicle for the peripheral expression and manipulation of information, drifting in and out of focus according to the ebb and flow of user activity.

2.1.2 TUI Precursors: Graspable User Interfaces

Compared with the vision of ubiquitous computing as the "third wave in computing" (Mark Weiser), the concept of "graspable" user interfaces was relatively narrow in scope. However, this focus on a particular aspect of the ubiquitous computing research agenda, drawing on previous work on bimanual styles of computer input (see Section 4.6.1), allowed George Fitzmaurice, Bill Buxton and Hiroshi Ishii to develop the hugely influential Bricks system (Fitzmaurice et al., 1995) and establish the conceptual foundations of Graspable UIs. These build on the notion of *graspable function handles*, in which the physical input device is tightly coupled with an associated virtual function. Such graspable function handles have five characteristic properties; they are (Fitzmaurice, 1996):

1. *space-multiplexed*, each virtual function having a single dedicated handle;
2. *strongly specialised*, each handle tailored to suit its single virtual function;
3. *concurrent*, multiple handles acting as parallel points of user control;
4. *spatially aware*, handles having positions within a spatial frame of reference;
5. *spatially reconfigurable*, handles being free to move within that reference frame.

Whilst these properties still hold in terms of current understanding of TUIs, the idea of graspable function handles is fundamentally 'pre-tangible'; it is concerned with *physical extensions of graphical user interface elements*, rather than predominantly physical interfaces consisting of meaningful physical representations. By the time the Graspable UI paradigm was evaluated by Fitzmaurice and Buxton (1997) (see Section 2.4), the concept of graspable function handles had been refined into a more general conceptual framework that attempted to unify the concept of graspable media with Weiser's concept of calm technology.

2.1.3 TUI Origins: Tangible Bits

The term "Tangible User Interface", and the "TUI" acronym, were both coined in Hiroshi Ishii and Brygg Ullmer's (1997) paper on *Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms*. Their aim was to "make computing truly ubiquitous and invisible" by taking advantage of "natural physical affordances (Norman, 1988) to achieve a heightened legibility and seamlessness of interaction between people and information"; this would be accomplished through TUIs that "augment the real physical world by coupling digital information to everyday physical objects and environments".

Weiser's influence is felt in the distinction between *graspable media* used to manipulate information at the centre of users' attention in the foreground of their activity, and *ambient media* used for passive awareness of information on the periphery of users' attention in the background of their activity. However, something is lost from Weiser's vision, in that the distinction between graspable and ambient media reifies the distinction between centre and periphery as categories of the world rather than states of the mind. Much TUI research has proceeded along one or the other of these parallel paths, with no real effort spent on developing TUIs that can drift in and out of our attention with the demands of real activities in real contexts.

The Tangible Bits vision can also be seen as somewhat of a false start, in that it was constrained by the old metaphors of graphical user interfaces. For each existing GUI element (windows, icons, menus, handles and widgets), an analogous TUI element was presented (lenses, phicons, trays, phandles and instruments respectively). Whilst such TUIs made use of meaningful physical representations, they were concerned with *physical instantiations of graphical user interface elements*, rather than the creation of new interaction metaphors and modes of representation. However, it could be argued that such an intermediate step was necessary in order to provide a conceptual stepping stone to a more ambitious model of TUIs free of unnecessary constraints. In any case, this next step was taken by the same authors three years later, in which they laid out a solid conceptual foundation that firmly established TUIs as a new and growing research area.

2.1.4 TUI Evolution: Emerging Frameworks

The paper on *Emerging Frameworks for Tangible User Interfaces* (Ullmer and Ishii, 2000) drew together much of the work on TUIs that had followed the Tangible Bits paper of 1997, and presented a clear conceptual structure and definition of TUIs that distinguished them from traditional graphical user interfaces. By reference to the Model-View-Controller (MVC) structure of GUIs, in which there is a clear demarcation between input and output, the structure of TUIs is presented as *Model-Control-Representation (physical + digital)* or *MCRpd*. The key characteristics of this interaction model are as follows (Ullmer and Ishii, 2001):

1. Physical representations (*rep-p*) are computationally coupled to underlying digital information (*model*);
2. Physical representations (*rep-p*) embody mechanisms for interactive control (*control*);
3. Physical representations (*rep-p*) are perceptually coupled to actively mediated digital representations (*rep-d*);
4. The physical state of the interface artefacts partially embodies the digital state of the system.

During the period between publication of this article in the *IBM Systems Journal* (Ullmer and Ishii, 2000) and its revised publication in the book *HCI for the New Millennium* (Ullmer and Ishii, 2001), "associative" TUIs – which "do not integrate the associations of multiple tangibles into larger-scale relationships" – were dropped as a category of TUI, since the authors were "less confident of the utility" of this kind of mapping. Instead, the additional constraint that the physical relations between tangibles should be detected and computationally interpreted – (4) above – was added to the TUI characteristics made explicit by the MCRpd model. This change was significant, because it clearly differentiated TUIs from other forms of physical user interface that were not based on systems of interacting objects. However, I will argue in this dissertation that the marginalisation of associative 'TUIs' fails to consider the value of user freedom in the creation of physical structures meaningful to themselves, but not necessarily to the system.

The MCRpd model now underpins almost all physical interfaces that employ systems of digitally augmented physical objects. The tight coupling between physical control and digital representation is the evolution of Fitzmaurice's graspable function handles, although the different forms of relationship between representations and the underlying model point to more varied semantic functions. In particular, Holmquist et al. (1999) refer to three main types of representation in TUIs: *tokens* are "objects that physically resemble the information they represent in some way"; *tools* are "representations of computational functions";

and *containers* are generic objects for the transient storage, manipulation and distribution of digital information.

The key breakthrough of the paper was that it presented TUIs in a manner emphasising the *physical embodiment of any kind of digital information*. It provided a platform for TUI research into systems that could not adequately be represented by instantiations of GUI metaphors, and provided an extensive literature review of the nascent research area.

If the MCRpd view of TUIs is to be criticised, then it is for its computational approach to TUI design. Whilst it has provided a platform for understanding issues of representation and control that bridge the physical and digital domains, as well as delineating different categories of TUI, it does not relate these concerns back to the contexts in which TUIs are to be deployed. This is of little relevance for demonstrations of design concepts, but when designing TUIs for real use, an appreciation of the deployment context is crucial for the development of an appropriate design. The *Emerging Frameworks* paper goes as far as saying *what* theoretical perspectives might influence TUI design, but not so far as saying *how* they should.

2.1.5 TUI Evolution: Embodied Interaction

Advice for TUI designers about how to create contextually appropriate designs can be found in Paul Dourish's (2001) book entitled *Where the action is: the foundations of embodied interaction* – an exploration of sociological and philosophical stances and theories relevant to the study of Human-Computer Interaction. The book concludes with six design principles that operationalise these diverse critical perspectives, all related to *embodied interaction*, or the "creation, manipulation, and sharing of meaning through engaged interaction with artefacts" (ibid.):

1. *computation is a medium;*
2. *meaning arises on multiple levels;*
3. *users, not designers, create and communicate meaning;*
4. *users, not designers, manage coupling;*
5. *embodied technologies participate in the world they represent;*
6. *embodied interaction turns action into meaning.*

The first pair of design principles relate to the fact that interactive technologies have meaning not just through their functionality, but through their symbolic social value and role within systems of practice. The nature of such technologies as interactive media jointly embodies the production of meaning by actions *on* the technology, as well as the sharing of meaning between people by actions *through* the technology. The implication for TUI design is that the characteristics of the interactive medium presented by the technology should take into account its future relationships to current technologies and practices.

The second pair of design principles build on the first pair, in that they relate to the ways in which users will appropriate interactive technologies to suit their individual and collective needs, within a complex web of cognitive and social activities. If design is carried out with reference to existing forms of contextual meaning, potential appropriations can be considered and incorporated into the functionality and presentation of the technology. The implication for TUI design is that the ways in which users collectively engage with the technology over time and space is determined by the meaning they have created from it, as well as the intended systems of meaning with which the technology was delivered. The role of the TUI designer is to find the correct balance between the under-specification of meaning (with consequences for usefulness) and its over-specification (with consequences for usability).

The final pair of design principles are most general, concerning the relationships between a technology and its environment. It builds on the second pair of principles, in that whilst the design of a technology affects its potential usage, adopting and adapting the technology by appropriation has a reciprocal effect in shaping the intersubjective systems of meaning that define the usage context. The implication for TUI design is that the primary concern should be how actions are rendered meaningful in social contexts, with the choice of media following on as a consequence of how this can best be achieved.

This 'primacy of action' marks another evolution of the TUI concept, by focusing on the *socio-digital effects of physical actions*. Rather than a computationally-motivated approach to TUI design, focusing on interfaces and information, this perspective encourages thinking in contextual terms about actions and meaning. However, while this contextually-motivated design approach to aims to produce more contextually appropriate interfaces, its high-level principles are less accessible to designers than low-level considerations of physical forms and digital mappings.

2.1.6 TUI Directions: Focus on Activity

The apparent tension between the computational and the contextual approaches to TUI design can be resolved by a focus on *activity*, which necessarily entails a conceptual linkage between the objects of activity (information via an interface) and its conditions and consequences (goals and meaning). Activities can be seen as proxies for the needs of specific users in specific contexts that allow empirically grounded theories of TUIs to be generalised to similar interfaces, places and people.

One of the most popular forms of activity that TUIs have been designed to support since their inception has been *co-located collaboration*. Recent work by Hornecker and Buur (2006) has resulted in a tentative framework for physical space and social interaction, highlighting the primary concerns in the design space of "collaboration-sensitive" TUIs for "socially-organized settings" (Williams et al., 2005). The paper also highlights the limitations and narrow focus of the traditional conception of TUIs, and complements the data-oriented view with perspectives taken from other disciplines: an "expressive movement"-centred view taken from Industrial and Product Design, and a space-centred view derived from the Arts and Architecture.

The framework itself consists of four themes, each elaborated by a set of concepts that aid in the understanding and application of those themes in TUI design; it presents a "deliberately non-restrictive view" of tangible interaction that enables "systematic shifts of focus", rather than strict design prescription. The four themes and associated concepts are as follows (Hornecker and Buur, 2006):

1. *Tangible Manipulation* refers to "the reliance on material representations typical for tangible interaction", through the concepts of *haptic direct manipulation*, *lightweight interaction* and *isomorph effects*.
2. *Spatial Interaction* refers to "how tangible interaction is embedded within space and occurs in space", through the concepts of *inhabited space*, *configurable materials*, *non-fragmented visibility*, *full-body interaction* and *performative action*.
3. *Embodied Facilitation* refers to "how configurations of objects and space affect social interaction by subtly directing behaviour", through the concepts of *embodied constraints*, *multiple access points* and *tailored representations*.
4. *Expressive Representation* refers to "the legibility and significance of material and digital representations", through the concepts of *representational significance*, *externalization* and *perceived coupling*.

These four themes succeed in capturing many of the aspects of tangibility. This very broad view of tangible interaction means that the framework is applicable to interfaces exhibiting any form of tangibility, but especially those involving co-located, collaborative, focal interaction. Further work has been carried out by Hornecker et al. (2007) into "how shareability comes about", explained through the concepts of *entry points*, which "denote design characteristics that invite people into engagement with a group activity and entice them to interact", and *access points*, which "denote characteristics that enable the user to actually interact and join a group's activity". Considering these concepts during the design of "shareables" encourages a focus on the *activity* rather than the *technology*.

I believe that a focus on activities will characterise the next wave of TUI research, as the concept that unifies the nature of interaction (e.g. tangible manipulation or spatial interaction) with its goals (e.g. embodied facilitation or expressive representation). Future chapters of this dissertation will present a similar analysis to Hornecker and Buur (2006) but in an almost diametrically opposed activity space: that of locally-distributed, coordinated peripheral-interaction. In the remainder of this chapter, I will present an overview of TUI research to date as part of the motivation for this new research direction.

2.2 TUI Design

TUI design conventionally follows the three stage process of motivate–prototype–analyse. These stages will now be discussed in turn.

2.2.1 Motivate

The way in which a TUI is motivated has a significant bearing on its ultimate design; in general, design can be regarded as being opportunity-driven, with the source of opportunity determining the nature of the ensuing design process. The two primary sources of opportunity are the realisation of a technology and the identification of a user problem – these are mutually beneficial, in that technologies provide the basis for currently feasible solutions, and the search for problems suggests new ideas for future technologies. However, the identification of user problems through contextual inquiry (incorporated by Beyer and Holtzblatt (1998) as a component of their Contextual Design process) has greater potential to lead to creative insight than the development of a technology, the demonstration of which typically involves reference to a motivating metaphor, rather than a motivating problem. With respect to the "Demo or Die" culture of the MIT Media Lab (of which Ishii's Tangible Media Group is part), Blackwell (2006) states that "the demo culture of research presumes that a realized design embodies metaphorical and other theories". In the same paper, Blackwell reports the participants of a 2003 workshop were "lamenting the loss of 'magic' in the over-literal desktop", and makes the following case against the "desktop metaphor" (ibid):

> *The desktop metaphor is not only mundane (an office rather than a laboratory), but moribund. It has been fully integrated into the everyday language of the UI, with no further potency as a generative creative image, one of those everyday dead metaphors that Nietzsche says are "worn by frequent use and have lost all sensuous vigour".*

The lesson for TUI design is that there *need not* be a motivating metaphor. This is not to say that metaphors aren't useful – they can be an invaluable tool for the presentation of design concepts which borrow from other domains. However, by searching for 'magic' metaphors that fully explain the workings of our designs, we may be led towards the creation of over-literal TUIs in the same way that the desktop metaphor has led to the creation of over-literal GUIs. The conclusion is that an interface metaphor should be recognised for what it is: a partial explanation of the partial workings of an interface; a mnemonic device that helps users to remember and communicate the ways in which an unfamiliar interface operates, by relating it to interfaces or domains with which users are familiar. A problem arises, however, when the perceived need for familiarity places artificial constraints on the types of interface investigated in the design process. Writing about the applicability of real-world behaviours in virtual environments, Hinckley et al. (1997) say that:

> *Virtual environments offer opportunities to violate the limitations of physical reality, and one only needs to mimic those qualities of physical reality which facilitate skill transfer or which form essential perceptual cues for the human participant to perform his or her tasks.*

TUIs also offer opportunities to enhance our interaction experiences by going beyond the limitations of physical reality, through the design of physical structures whose salient perceptual cues are digitally augmented in a way that directly facilitates skill transfer from everyday physical experiences to the underlying digital application. This is a very different kind of design to the identification and mimicry of a literal domain, and requires creative techniques for the exploration of ambiguous, indeterminate and semi-formulated physical design concepts.

2.2.2 Prototype

A valuable tool in the HCI designer's toolkit, capable of addressing both dead metaphors and technology-driven design, is *low-fidelity prototyping*. In the world of professional product design, the role of sketching out design concepts to see where they might lead is known as "conversation with the materials", a central aspect of the reflective design practitioner's experience (Schön, 1987). Our paper on *Using Solid Diagrams for Tangible Interface Prototyping* (Blackwell et al., 2005) is an attempt to introduce such practices to the design of TUIs. We describe an approach to rapid, low-fidelity *tangible prototyping*, in which three-dimensional sketching materials are used for creative exploration of the physical design space. Such materials include foam boards, plasticine modelling clay, wooden spatulas, pipe cleaners, string, stickers, paint, glue sticks, and a variety of other fasteners and decorative attachments (see Figure 2.1). These sketching tools allow the creation of what we call *manipulable solid diagrams*: physical objects that can be analysed as a means of representing and controlling an underlying information-structure (Edge and Blackwell, 2006). In general, two-dimensional diagrammatic representations are "indexed by location in a plane" and "typically display information that is only implicit in sentential representations and that therefore has to be computed, sometimes at great cost, to make it explicit for use" (Larkin and Simon, 1987). The challenge for the design of manipulable solid diagrams is to take advantage of their interactive, material properties whilst retaining the characteristic inferential properties of diagrammatic representations. Such creations can be subjected to evaluation by abstract analysis (of the kind discussed in the next chapter), with the results feeding back into the next generation of prototypes within an iterative design process. Since most potential users of TUIs are unlikely to be familiar with the concept of 'tangible' interfaces and post-WIMP interfaces in general, the use of tangible prototyping as a participatory design technique for envisioning future technologies is a valuable tool in tangible interaction design. In our paper (Blackwell et al., 2005), we describe the format of a typical workshop session in which our solid diagram prototyping is conducted with such potential users.

In contrast to our "3D sketching" approach to TUI design, the *simplicity design exercise* of Chang et al. (2007) is a lightweight approach based on traditional pencil and paper sketching. The idea is to "explore the communication capacity for interaction components" by redesigning existing devices according to two constraints: choose a minimal number of interaction mechanisms (e.g. LEDs, buttons) with at most one of each type, and explore the combined expressiveness of these mechanisms in the form of a state diagram. Whilst this is a useful exercise, it is not intended as a real design tool. In contrast, *Papier-Mâché* (Klemmer et al., 2004) is intended for real use as a high-fidelity prototyping toolkit, suitable for the construction of TUIs based on computer vision, electronic tags, and bar-codes. The API of the toolkit abstracts away from the details of individual sensing technologies and presents a high-level event model that aims to encourage both rapid prototyping and technology portability.

Figure 2.1: Modelling Materials for "3D Sketching" of TUI designs

2.2.3 Analyse

After constructing a low-fidelity TUI prototype, it is desirable to understand how it compares to existing TUIs and how it will be interpreted as meaningful by its users. Two related ways of describing a TUI are to refer to its *structural form* and its *style of mapping* (Ullmer and Ishii, 2001). The structural form of a TUI defines how meaningful physical structures are created: by the position and orientation of tangibles on an *interactive surface* that provides digital feedback; by the *constructive assembly* of tangibles into unitary structures; or by the relative configuration of tangible *tokens* within a system of physical *constraints*, known as the *token+constraint* structural form. The complementary analysis of a TUI's *style of mapping* defines the way in which physical configurations are interpreted as meaningful: through the *spatial* location of tangibles within some absolute frame of reference; through the *relational* arrangement of tangibles relative to one another; or through the *constructive* connection of tangibles into meaningful structures.

Whilst the notion of token+constraint interfaces was introduced by Ullmer in his dissertation on *Tangible interfaces for manipulating aggregates of digital information* (Ullmer, 2002), it has since been refined and generalised as a model of all TUIs, regardless of their structural form, under the acronym "TAC" (Calvillo-Gámez et al. 2003, Shaer et al. 2004). The "TAC paradigm" is a way of analysing TUIs by systematically describing all of the relations between their physical and digital parts. In the resulting description language, the term "pyfo" is used to denote a physical object that takes part in a tangible interface, with each pyfo being a token, a constraint, or both. A token is then a pyfo that represents a variable, be it some digital information or a computational function, and a constraint is a pyfo that limits the behaviour of the token with which it is associated. A TAC is therefore the (perhaps temporary) relationship between a token, its variable, and one or more constraints. The five key properties of TUIs from the perspective of the TAC paradigm are as follows:

1. *Couple*: a pyfo must be coupled with a variable in order to be considered a token.
2. *Relative definition*: each pyfo may be defined as a token, a constraint, or both.
3. *Association*: a new TAC is created when a token is physically associated with a constraint.
4. *Computational interpretation*: the physical manipulation of a TAC has computational interpretation.
5. *Manipulation*: each TAC can be manipulated discretely, continuously, or in both ways. The physical manipulation of a token is afforded by the physical properties of its constraints.

The primary purpose of the TAC paradigm is as a descriptive tool, although a high-level visual language for specifying TUIs and other "reality-based" interfaces has been attempted (Shaer and Jacob, 2006). This represents a move closer towards the goal expressed in Calvillo-Gámez et al. (2003) of identifying TUI design patterns – elegant solutions to commonly occurring problems in a design context – and presenting them as a taxonomy.

A more abstract classification of tangible interfaces is presented by Dourish (2001), segmenting the design space along two dimensions: actions–objects, i.e. to what extent the elements of the interface represent actions or objects; and iconic–symbolic, i.e. to what extent the physical representation has similarity with the digital thing it represents. A similarly abstract classification is given by Fishkin (2004), who classifies tangible interfaces according to the two dimensions of embodiment (attention to the object of manipulation) and metaphor (analogy to real-world actions). Whilst all of these approaches are useful for analysing designs that already exist, they are not generative, in that they do not give explicit design guidance. The next chapter of this dissertation will address this gap in the literature by presenting a wide range of analytic tools and describing how they can be combined within an *analytic design process* for TUIs. The remainder of this chapter will present an overview of TUI systems and their evaluation

2.3 TUI Systems

My presentation of TUI systems is based on a six-way categorisation of the TUI design space. The primary distinction is made based on the overall purpose of the TUI: as a *direct-manipulation model*, in which the effects of actions are experienced locally in both time and space; as a *behavioural specification*, in which the effects of actions are experienced at some future time; or as an *communication channel*, in which the effects of actions are experienced immediately at some remote location. These can be seen as the implicit metaphors of TUI design, which have emerged over time as alternatives to the traditional metaphors of graphical interfaces. A secondary distinction is then made within each metaphor, between the different modes of representation outlined by Dourish (2001): *iconic representations* are those tangibles that represent domain objects through similarity, whilst *symbolic representations* are those tangibles that represent by arbitrary association.

Photographs of selected TUI systems from the following sections are shown in Figure 2.2.

2.3.1 Tangibles as Direct-Manipulation Models

We have noted, in (Blackwell et al., 2007), that “most [TUI systems] focus on the immediate effect of communication to provoke system action or change of state”. This direct-manipulation metaphor is often implicit in arguments about TUIs leveraging our natural abilities and familiarity with the physical world, since it follows ’natural’ expectations of causation.

Iconic Models

TUIs based on iconic models mimic physical situations for a variety of purposes. A popular foundation for such interfaces is an iconic map-like surface (which itself is an iconic representation of some external territory). The first instance of this was the Tangible Geospace application running on the metaDESK as described in Ishii and Ullmer (1997), in which the relative configuration of two building “phicons” is used to position, scale and rotate a projected map of the MIT campus. In addition to landscape visualisation, simulation is also possible with such interfaces. The MapTable (Reitmayr et al., 2005) goes even further into the physical domain by digitally augmenting paper maps in real time, for the purpose of flood defence planning and disaster management. In the MapTable interface, the paper maps retain the physical affordances of paper and can be annotated, overlaid, reoriented and reordered. Moving tangible “plates” to different hot spots on the augmented map results in the projection of live video feeds from those locations (see photograph (a) in Figure 2.2); similarly, placing tangible “frames” anywhere on a map results in additional local information being displayed within their boundaries. The multimodal augmentation of wall-based maps has also been explored in the Rasa system (McGee and Cohen, 2001), in which the locations, annotations and voice descriptions of sticky notes are detected, and their subsequent rearrangement used for real-time command and control.

In the systems described so far, the primary source of iconic similarity with the real world is the underlying surface. An even more iconic ’surface’ is the Illuminating Clay system (Piper et al., 2002), in which digital feedback is projected onto hand-sculpted clay landscapes (see photograph (b) in Figure 2.2). A well-known TUI that demonstrates the alternative – iconic objects in meaningful arrangements on a meaningless surface – is the Urp system (Underkoffler and Ishii, 1999) for urban planning. Here, models of buildings are arranged as would be the case in reality, with the surface used to simulate the effects of wind and time of day. Similar use of a plain surface for simulation is found in the Illuminating Light system (Underkoffler and Ishii, 1998) for the experimental arrangements of mock optical components in the production of holographs.

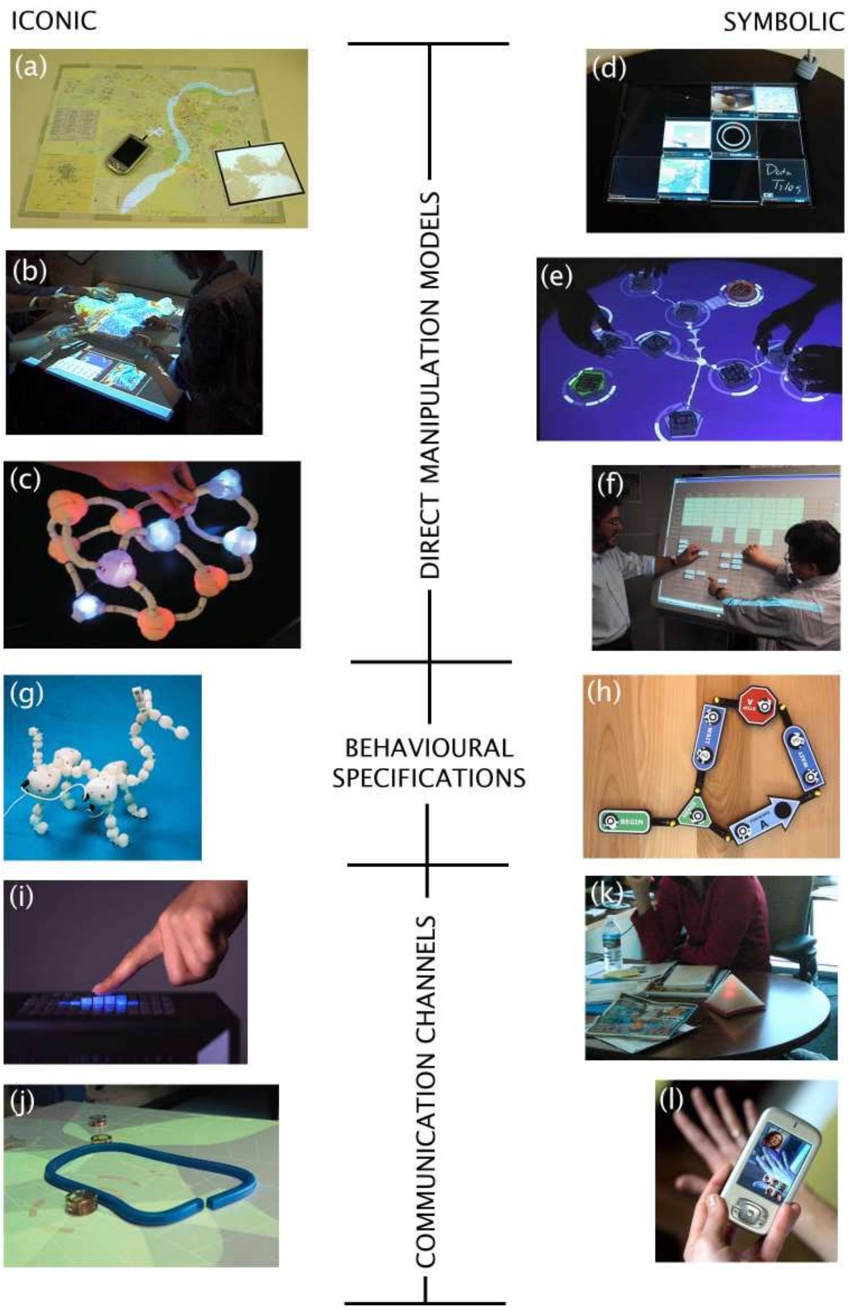

Figure 2.2: Photographs of Selected TUI Systems: (a) MapTable (Reitmayr et al., 2005); (b) Illuminating Clay (Piper et al., 2002); (c) Senspectra (LeClerc et al., 2007); (d) DataTiles (Rekimoto et al., 2001); (e) reacTable (Jordà et al., 2007); (f) Senseboard (Patten et al., 2001); (g) Topobo (Raffle et al., 2004); (h) Quetzel (Horn and Jacob, 2007); (i) Lumen (Poupyrev et al., 2004); (j) PICO (Patten and Ishii, 2007); (k) Nimio (Brewer et al., 2007); (l) Telebeads (Labrune and Mackay, 2006).

Iconicity can come from forms other than the spatial arrangement of objects. The Senspectra system (LeClerc et al., 2007) makes visible the invisible stress forces within physical structures (see photograph (c) in Figure 2.2), whilst the Passive Props system (Hinckley et al., 1994) for neurosurgical visualisation makes visible the internal cross sections of brain, through the placement of a two-dimensional surface against a physical model of the head. An alternative to using digital augmentation for making visible the otherwise invisible, is to subvert reality and make visible the imaginary. A good example of this is the I/O Brush (Ryokai et al., 2004), which mimics the physical conventions of painting yet allows the stroking of textures 'picked-up' from the environment. These textures can be static, acquired from photographs taken by the "brush", but also dynamic, acquired from video recordings taken similarly. In a related but different kind of subversion, the Siftables prototype (Merrill et al., 2007) uses compact graphical displays that iconically display their contained media. The system is an example of a Sensor Network User Interface (SNUI), in which individual displays can determine their location and motion to such an extent that they can be manipulated individually and in groups, responding to physical gestures and supporting literal 'sifting' through the underlying media.

Symbolic Models

Model-based TUIs need not be limited to existing physical scenarios, however: symbolism can be used to represent abstract concepts and the relationships between them. There are three general themes within this category: objects used to support the convenient organisation and access of information; objects used for real-time performance due to their interaction directness; and objects used for joint problem solving due to their visibility and malleability.

Symbolic objects for convenience: media management

One of the best examples of tangible media management, often held up as a design exemplar but before the coining of the term "TUI", is Durrell Bishop's Marble Answering Machine (Polynor, 1995). In this design, physical marbles are used for the representation and control of voice messages. Similar symbolic representations of intangible media are found in a variety of systems, such as the early mediaBlocks TUI (Ullmer et al., 1998) in which cube-shaped blocks are used for the transportation of online media and its control within linear rack constraints. More complex relationships between intangible applications are possible in the DataTiles system of Rekimoto et al. (2001), based on transparent tiles manipulated within two-dimensional grid-based trays (see photograph (d) in Figure 2.2). Such tiles are orthogonally coupled for the purposes of binding parameters to applications, recording interactions with applications, and broadcasting applications to other tile trays. A more lightweight approach to media management is seen in the Tangible Tiles system (Waldner et al., 2006) in which transparent tiles are used to interact with digital media and transport it between multiple display surfaces.

Whereas all of these approaches have associated media with modular physical objects, there is an alternative approach. The Clutter Bowl project (Taylor et al., 2007a) uses the well known understanding of bowls as objects of convenient containment within the home, and symbolically associates the placement of digital media devices in the bowl with the projection of their media content around the bowl's inner surface.

Symbolic objects for convenience: control management

The VoodooIO system (Villar and Gellersen, 2007) aims to support individual computer users at their desks by providing "network substrate" materials that support the "graspable function handle" concept of Fitzmaurice – each control has connecting pins that perforate this substrate, allowing complete freedom of control layout. The "malleable control structure" in terms of control inclusion and spatial layout is coupled with virtual commands through a process called "softwiring", in which physical controls can either be statically pre-associated with particular virtual commands, or dynamically configured via the sequence of: 1) select virtual command; 2) add new control to substrate. This generic workstation support system has been instantiated and evaluated as an extension of the Adobe Flash application for dynamic media creation (Spiessl et al., 2007), and was also influenced by the similar but specialised Pin & Play & Perform system (Villar et al., 2004), whose applicability was confined to the control of MIDI parameters. The Amphibian project (Carvey et al., 2006) is a novel control management system also used for general support of workstation interactions, but based on the appropriation of existing physical objects as weight-indexed representations of computational functions. Any object can be associated with a computational function – keypress combination or macro – by registering its weight on a set of sensitive USB scales and mapping that weight to the desired function. The entire desktop thus becomes an environment of control convenience, with functions activated on demand by placing the desired object back onto the USB scales.

Symbolic objects for real-time performance: musical improvisation

The reacTIVision system is one of the most polished TUI infrastructures to date, and Kaltenbrunner and Bencina (2007) present a guide to the construction of an interactive-surface using the open-source components made available by previous contributors to the project. The resulting surface requires rear-mounting of both a projector and a webcam modified to work in the IR range. The project uses a novel genetic algorithm-based approach to fiducial marker[1] generation – 90 such "Amoeba" fiducials are distributed with the system for object tagging, resulting in sub-pixel localisation accuracy as well as orientation. The most widely recognised use of the reacTable is for live music performance (see photograph (e) in Figure 2.2), discussed in Jordà et al. (2007). The AudioPad system of Patten et al. (2002) is another surface-based TUI for live music performance using a bimanual style of parameter control, whilst Block Jam (Newton-Dunn et al., 2003) is a block-based, two-dimensional constructive assembly tightly coupled to a musical sequencer. Finally, the Music Cube (Alonso and Keyson, 2005) is a simple tangible appliance for the performance of music playback.

Symbolic objects for real-time performance: narrative

[1] A "fiducial" is used as a fixed standard of reference for comparison or measurement (WordNet 3.0); in computer-vision based systems, "fiducial markers" are printed codes that can be attached to objects to identify their position and orientation in space. Two different coding schemes are described in López de Ipiña et al. (2002) and Kaltenbrunner and Bencina (2007).

Symbolic objects have also been applied in narrative systems, where the configuration or sequential selection of tangibles is used to support the evolution of stories, arguments or presentations. The early Triangles system (Gorbet et al., 1998) "provides a physical embodiment of a digital information topography" that can be used for non-linear storytelling. In the system, connections between modular triangles are interpreted, triggering audible story events based on both the comic book style graphics affixed to the triangle surfaces, and on the previous event history. Whereas this system is based on physical connection, the Tangible Viewpoints system (Mazalek et al., 2002) uses the number, position and adjacency of character "pawns" to control the unfolding narrative based on the story segments displayed in the coloured "auras" surrounding each character pawn. The StoryGrid project (Moher et al., 2005) bears a surface similarity to the Tangible Viewpoints system, but is aimed at group interpretation of depicted scenarios rather than narrative generation. Even more focused on group narrative support is the classroom-based Webkit system (Stringer et al., 2004) in which transparent tiles are associated with online media. These tiles have two usage phases: firstly as a means of collaborative argument construction within linear rack-like constraints, and secondly as a means of information access during presentation of that argument to other children. Finally, the Palette system (Nelson et al., 1999) is a tangible representation of presentation slides that can be dynamically re-ordered and revisited during the course of a presentation.

Symbolic objects for real-time performance: other tasks

Whilst real-time narrative and musical improvisation are the primary application areas of symbolic performance, some task domains require similar levels of immediate access and direct manipulation. Two of the better known systems of this kind were developed in parallel at Interval Research: the LogJam system (Cohen et al., 1999) for group video logging, and the ToonTown system (Singer et al., 1999) for controlling multi-user presence within an audio space. Both of these systems require rapid, task-focused interactions when user attention is divided between the primary task of attending to video or voice, and the auxiliary task of responding to events with the toggling of log codes or the reconfiguration of the audio space respectively.

Symbolic objects for problem solving in abstract domains

A final use for symbolic objects is not for convenience or performance, but for problem solving and design. Here, the tangible elements participate in abstract relations that can be reasoned about and manipulated by multiple users at the same time. The Senseboard (Patten et al., 2001) platform, and subsequent applications such as IP Network Simulation (Kobayashi et al., 2003), are typical of this TUI category. Such abstract interfaces are often driven by new sensing technologies; in the case of the Sensetable, this was the creation of a sensing mechanism capable of robustly identifying and tracking the positions and orientations of multiple symbolic "pucks" containing a simple type of embedded RF (Radio Frequency) tag consisting of an inductive wire coil and a capacitor. Their different resonant frequencies, combined with specific antenna geometries, enable tracking of up to nine pucks with both high resolution and low latency. The Sensetable pucks were also one of the first tangibles to contain embedded physical controls, in the form of parameter dials. A recent evolution in interactive surface technology is ThinSight (Hodges et al., 2007), which uses IR sensing through thin form-factor displays to effectively unify tangible and multi-touch interaction. Whilst the sensing resolution is currently too low for barcode-like identification of tangibles, this is likely to become possible in the near future.

In contrast to this current work, one of the earliest interactive surfaces was provided by the BUILD-IT system (Rauterberg et al., 1998), which relies on similar symbolic association between physical blocks and domain objects. A system different to ThinSight and BUILD-IT in both surface type and structural form is the Senseboard (Jacob et al., 2002) – a wall-mounted surface composed of a regular two-dimensional grid, used to constrain the attachment of rectangular blocks to discrete rectilinear arrangements (see photograph (f) in Figure 2.2). This structure is especially suited to the representation of object grouping and order, and was used for this purpose in its evaluation as an application for the collaborative organisation of academic papers into conference sessions, and the ordering of the paper presentations within the created sessions. The Designers' Outpost (Klemmer et al., 2001) is another wall-based system, this time for the collaborative structuring of Websites using sticky notes and an interactive whiteboard, whilst the FlowBlocks system (Zuckerman et al., 2005) is altogether different, being a classroom-based constructive assembly for the understanding of probability distributions built from symbolic path blocks.

2.3.2 Tangibles as Behavioural Specifications

Following on in Blackwell et al. (2007) from our argument that most TUI systems focus on the immediate effects of interaction, we point out that "a further alternative is for the user to specify the structure of the required behaviour, rather than directly specifying the required actions". Such interfaces no longer provide direct manipulation, since the effects of changes are only visible once the 'specification' is activated, and even then only to the partial extent that the prevailing input conditions determine one set of behaviours out of potentially many. We have referred to such TUIs as *manipulable solid diagrams* (Edge and Blackwell, 2006) to emphasise their abstract representational structure and its inherent malleability.

Iconic Specifications

The iconic specification of behaviour is akin to programming by demonstration, in that the resulting behaviours (the output effects) resemble their specification (the input actions), albeit at some future point in time. This characteristic is the defining feature of the Topobo system (Raffle et al., 2004) – a modular construction kit with kinetic memory, used to create both biomorphic and mechanical constructions (see photograph (g) in Figure 2.2). Manipulations of active rotary components can be recorded as specifications of behaviour, and 'played back' in a repeating loop. Using such actions, two users can specify the gait of a quadruped creature by simultaneously manipulating all of its four legs. The use of special *queen* components introduces an additional abstraction facility into the specification process, since the manipulations recorded by a queen will be replayed by all active components transitively connected to that queen. Different types of queen can be used to create different types of replication: regular queens result in exact, simultaneous copies of motion; "decay" queens result in simultaneous motion but with a reduction of amplitude that scales with distance from the queen; "time delay" queens result in the amplitude of motion being replicated but with an increase in time delay that scales with distance from the queen; whilst "faster/slower" queens results in harmonic resonance patterns within connected chains of active components. This notion of "queen" components has since been updated by the concept of *backpacks*, presented as "tangible modulators for kinetic behaviour" (Raffle et al., 2006). These allow the recorded motion of active components to be individually modified before playback occurs, with specification of three possible types of influence: local influence affects the active component to which it is attached (the default); global influence identically affects all active components in the structure (by pressing the button on that backpack); and distributed influence results from the attachment of a backpack to a queen. Backpacks therefore allow a distinction to be drawn between the replication of behaviour by queens, and the modification of behaviour by backpacks. Note that the use of these abstractions is not necessary, but their introduction reduces the iconicity of behaviour specification. A comparatively simple tangible appliance that is purely iconic is the floor-based Curlybot (Frei et al., 2000), an "autonomous two-wheeled vehicle" that can record and play back its motion on an underlying surface.

Symbolic Specifications

Symbolic specification in space

The iconic nature of tangible "programming by demonstration" discussed in the previous section is in contrast to the highly symbolic nature of programming in general. Two of the earliest 'tangible' interfaces for symbolic programming were AlgoBlock (Suzuki and Kato, 1993) and the Slot Machine (Perlman, 1976) – both providing physical programming languages for children to solve problems in a subset of Logo. In the AlgoBlock system, programs were linear connections of blocks containing embedded knobs and levers for the specification of command parameters, whilst the Slot Machine was composed of three coloured racks into which cards representing Logo commands could be plugged, with parameters dealt with by stacking parameter cards on top of the larger command cards. In both of these systems, the physical configuration of objects specifies the behaviour of a Logo program. Modern incarnations of this concept are found in the Tern system (Horn and Jacob, 2007), in which commands for a virtual robot are represented by wooden blocks adapted from an existing children's railway construction kit, connected by jigsaw-like constraints. Quetzel (ibid.) is another tangible programming language by the same authors, which instead uses point connectors – freely rotatable in the plane of the underlying surface in order to create more complex control structures – to specify the behaviour of a LEGO Mindstorms robot (see photograph (h) in Figure 2.2). Another system using jigsaw-like connection, but concerning behavioural specification in a different domain, is the Tangible Video Editor (TVE) system (Zigelbaum et al., 2007). This uses a jigsaw metaphor to constrain the linear spatial construction of temporally-ordered video segments, with physical icons of fade, rotate and minimise transitions integrating in similar jigsaw-like fashion along the top edge of the resulting structure. Tangible video editing has also been addressed by Taylor et al. (2007b), in which magnetic rather than jigsaw connection is used to allow dynamic construction of video sequences and transitions on the Microsoft Surface[2].

Symbolic specification over time

A contrasting approach to tangible programming is seen in the MediaCubes system (Blackwell and Hague, 2001), which aims to extend the remote control metaphor for the networked home. The result is that programs are created by the reconfiguration of cubes over time, rather than space. Each cube has a single button that can be used to invoke some fixed function of an appliance associated with it. The cubes themselves can be temporarily associated with particular appliances, in order to refer to the function of that appliance. When the sides of two cubes are aligned, and the button on each cube pressed, a relationship is constructed between the operations represented by those two faces. A typical combination of MediaCubes might be to associate an "event" cube with a particular time, and a "channel" cube with video transmission from a satellite receiver to a DVD recorder. Aligning the "on" face of the event cube with the "video" face of the channel cube would then define a simple script to start recording at that time. The interface is tangible, even though the resulting programs are not.

2.3.3 Tangibles as Communication Channels

The TUIs discussed up to this point have all been concerned with the production of effects local to that instance of the TUI, and either immediately in time after the interface action (in what I have called *direct-manipulation models*), or at some future point in time depending on 'program' activation and input parameters (in what I have called *behavioural specifications*). A third alternative is for program effects to be observable immediately but at some other location, with TUIs being used as *communication channels*.

Iconic Channels

At the highest level, all human communication is symbolic since the interpretation of words, pictures and actions into mental structures is essentially a process of learned convention. However, we can still define an iconic channel as one in which locally performed actions are mirrored in remotely executed effects. One of the better known examples is the InTouch system (Brave et al., 1998), in which paired tangible devices, each constructed from three rolling cylinders arranged in parallel, can replicate their respective motion in the other by a process of "digitally mediated bilateral force-feedback". This is designed to support "haptic interpersonal communication", providing a means of shared expression through touch, in which distributed users can passively feel or actively oppose the actions of the other. Whilst the TUI iconically reproduces the input actions, these actions have no inherent semantics and so must be interpreted symbolically. A TUI in which the objects have greater potential for iconicity (see photograph (i) in Figure 2.2) is the Lumen Shape Displays

[2] http://www.microsoft.com/surface/

of Poupyrev et al. (2004). These support the iconic deformation and remote recreation of 3D physical shapes through 13×13 'pixel' arrays of actuated glowing rods, although the limited resolution constrains iconic expressions to such things as "smiley faces". Greater expressive potential is found in the PSyBench platform (Brave et al., 1998), which implements the concept of "synchronized distributed physical objects" whose movements are simultaneously replicated in distributed instances of the interface. Whilst the prototype of this concept was built using motorised chessboard technology, the Actuated Workbench (Pangaro et al., 2002) extends the Sensetable concept to incorporate electromagnetic actuation of pucks (with the consequence that pucks require computer-vision rather than RF tracking to avoid electromagnetic interference). Another approach to the synchronisation of remote surfaces is seen in the Remote Active Tangible Interactions (RATI) system (Richter et al., 2007), which uses surface-based robots in conjunction with rear projection. A turn-taking interaction style is used to deal with the problems of distributed users attempting to simultaneously move the 'same' robotic object. The choice of robotic or electromagnetic actuation appears to manifest a trade-off between concurrent and scaleable interaction, although this is not the most problematic aspect of synchronised TUIs. They must operate with a fixed number of tangibles of known types in a 'closed' system, negotiated a priori and effectively limiting what can be discussed in a manner that precludes the addition or removal of other tangibles to or from the surface. A partial solution is to employ an 'offline' area, with the placement or movement of tangibles into this area logically equivalent to removing that tangible from the application space. However, the problem still remains that the types and numbers of tangibles need to be negotiated between distributed users beforehand in order to function correctly, and that users are prevented from doing what they might naturally otherwise do with physical objects: reorient them to read their 'label', pick them up to annotate them, carry them in their hand or pocket, put them in a meaningful place for future use, and so on. These problems are mitigated to some extent, however, when communication is not with a remote user, but with the system. The PICO system (Patten and Ishii, 2007) demonstrates how puck actuation can be used as real-time feedback from an "algorithmic partner", within an activity of collaborative computational optimisation (see photograph (j) in Figure 2.2). To assist with their task of laying out mobile phone towers, users can place any available physical objects on the board to constrain possible puck motion, in a manner seamlessly integrated into the system's model of solution space.

Symbolic Channels

Output effects need not always mirror input actions. In symbolic communication channels, the input actions (at one or more TUIs) are interpreted and executed as symbolic effects (at one or more TUIs). A good example of this is the Nimio system (Brewer et al., 2007), in which translucent, geometric-solid "playthings" symbolically respond to remote sound and touch: individual input actions are translated into collective output effects of glow pattern and colour (see photograph (k) in Figure 2.2). The authors refer to it as a member of a class of TUIs that provide an "effective mechanism for combining ambient displays with social connection through activity awareness". Whilst this system is designed to support N-way passive awareness, the LumiTouch appliance (Chang et al., 2001) "attempts to enhance the symbolic power of the picture frame, by providing a subtle real-time communication link". Presence of a person at the remote picture location is passively communicated by an ambiently glowing frame, whilst actively picking up the frame and squeezing it results in symbolic patterns of colour being transmitted based on squeeze location and duration. Finally, the Telebeads system (Labrune and Mackay, 2006) is based on "mobile mnemonic artefacts" that allow teenagers to link individuals or groups to particular items of BlueTooth augmented jewellery, using a PDA-based "toolglass" to reveal these associations (see photograph (l) in Figure 2.2). The paper introduces the concept of *interperception* as the "peripheral or even subliminal awareness we have of each other", and relates this subjective symbolism to the complementary social symbolism of *experience networks*; in the Telebeads example, an individual's jewellery items act as indices to other members of the experience network, used to initiate both digital communication with remote members of the network, and social conversation with those members who are co-located. Whilst not a tangible communication channel in the same sense as the other systems discussed here, the use of tangibles in a way that exploits our existing subjective and social interpretation of 'things' is an interesting area for future research.

2.4 TUI Evaluation

The ways in which TUIs are evaluated roughly fall into six categories – in terms of increasing formality, these are: casual observation by colleagues or visitors; presentation as interactive exhibits at shows or museums; observation-based user studies with mock tasks; observation-based user studies with real tasks; extended contextual deployment; and controlled experiment. The distinction between mock and real tasks is made according to whether the same or similar task would have been performed in the absence of researcher intervention. Many of the papers on TUI systems presented in the previous section do not refer to any kind of evaluation at all. For some systems, there are good reasons for this – for instance the Marble Answering Machine (Polynor, 1995) is a design rather a functioning prototype, the Sensetable (Patten et al., 2001) is a platform rather than an application, and the MediaCubes (Blackwell and Hague, 2001) are a programming language concept for the networked home of the future, rather than of the present. However, whilst proofs of concept are certainly useful as design inspiration, the design concepts alone are not proof of being useful. More thorough investigation of the quantifiable performance enhancements (probably from controlled experiments) and the qualitative benefits experienced (probably from contextual deployment) are necessary to provide better foundations for future work on TUIs. The various types of TUI evaluations conducted are now presented in decreasing order of formality.

2.4.1 Controlled Experiments

The first controlled experiment relevant to TUI research was a between-subjects comparison of Fitzmaurice's graspable function handles to a regular mouse-operated GUI application, in a fixed-duration target-tracking task (Fitzmaurice and Buxton, 1997). The results of the comparison showed a significant difference between the time-multiplexed generic-device condition, the space-multiplexed generic-device condition, and the space-multiplexed specialised-device condition, with the conditions involving space-multiplexing and specialised-devices each resulting in significant improvements, in the root-mean-square tracking error, over those conditions involving time-multiplexing and generic-devices. The conclusion is that specialised form factors contribute towards reduced switching costs relative to generic form factors, and that space-multiplexing contributes towards reduced switching costs relative to time-multiplexing.

The benefits of tangibility are not purely confined to those of information access and manipulation efficiency. Patten and Ishii (2000) use a between-subjects design to compare how people use space to organise identical tangible blocks and graphical icons respectively into meaningful structures, from which the meaning of individual elements can be recalled. The results were that TUI subjects significantly outperformed GUI subjects in recall accuracy, that only TUI subjects thought of the whole available space as meaningful (as opposed to simply imbuing relative token positions with meaning), and that those TUI subjects who used this "reference frame based positioning" performed significantly better than those who did not.

Of recent TUI systems, one of the only studies to derive useful results from controlled experiments is the evaluation of the PICO system (Patten and Ishii, 2007). This experiment used a between-subjects design in a cellular tower placement task with three conditions: PICO (with actuation and physical constraints), PICO′ (without actuation or constraints), and screen and mouse (no actuation, but virtual constraints). Results showed that the number of times each subject switched objects was significantly higher in the PICO condition than the PICO′ condition, which in turn was significantly higher than the screen condition. Subjects also used constraints significantly more often in the PICO condition than the screen condition. To make comparisons across conditions, a one factor ANOVA was conducted to see if there was a relationship between the tendency to switch objects and successfully completing the task by reaching the target coverage-score within the fixed time limit. The result was that in trials where subjects successfully completed the task, their tendency to switch between towers was significantly greater than in those trials where subjects failed to reach the target score. The conclusion is that for collaborative optimisation tasks, actuated tangibles can encourage behaviour that is both more exploratory and efficiently goal-directed compared to tangibles alone, which in turn outperform screen and mouse solutions.

2.4.2 Extended Contextual Deployment

The often fragile nature of TUI demonstration systems means that most are unsuitable for deployment and unsupervised use in real contexts. One system that has been evaluated in this way is the Nimio system (Brewer et al., 2007) for passive, symbolic awareness of group activity. The design and contextual deployment of the system were motivated by an earlier site study examining the role of physical artefacts in the spatial and social context of work, of the same people to whom the system was later deployed. The ambiguous nature of communication was intended to focus user attention on the technology, letting new modes of communication develop organically, rather than being forced. The study itself was based on observation and semi-structured ethnographic interview of the ten people – managers of a multi-disciplinary information technology research institute – to whom Nimio devices were deployed. The main contribution of the work was to highlight the tension between the desire to create *ambiguous* designs that would engage users but not interfere with existing social norms and practices, and the desire to create *legible* designs that could be 'read' by users and used as a basis for intentional meaning-creation based on collective interactions over time.

2.4.3 Studies Involving Real Tasks

The context of the classroom is one where TUIs can be applied to 'real' educational tasks more readily than, for example, 'real' work tasks in an organisation, due to the relative freedom with which teachers can choose to deliver the curriculum. Whereas persuading a manager to substitute a TUI for their email client or similar is likely to meet with resistance, substituting a tangible painting application for real paint might actually be preferable to the teacher and provide an alternative and engaging mode of learning for the children. This is the application area of the I/O Brush (Ryokai et al., 2004), evaluated in the classroom context by observing the interactions of pairs of children with the system. The main conclusion drawn from observation was that "although the outcome of [the children's] artwork was synthetic and digital, the process of their work involved searching for and interacting with many physical objects that are available and meaningful to them in their life. Through such exploration with familiar objects and constructing meanings through them, children learn to take control over underlying abstract concepts". A similar observation-based evaluation in the classroom context, albeit for older children, was a series of eight preliminary trials using incremental generations of the Webkit interface (Stringer et al., 2004) for structuring and presenting rhetoric arguments based on source material from the Web. Evaluations of the Topobo system (Raffle et al., 2004) and its Backpacks modifiers (Raffle et al., 2006) have similarly been evaluated as an educational tool for children at various educational levels. An example of a TUI being evaluated based on the performance of real tasks is the Logjam system (Cohen et al., 1999), in which a group of video loggers used the TUI to log eight 1-2 hour tapes, spending around fifteen to twenty-five hours to do so. Such evaluations using live data are rare in the TUI literature.

2.4.4 Studies Involving Mock Tasks

A 'safer' option than using TUIs to perform real tasks in context is to create mock tasks that can be performed by potential future users. The general format of these studies is: 1) system introduction; 2) observation/recording of task performance; 3) survey-based feedback. This procedure was followed in the published evaluations of the Tangible Video Editor (Zigelbaum et al., 2007), Tangible Tiles (Waldner et al., 2006), Designers' Outpost (Klemmer et al., 2001), VoodooFlash (Spiessl et al., 2007), Palette (Nelson et al., 1999), Music Cube (Alonso and Keyson, 2005), and RATI (Richter et al., 2007).

2.4.5 Interactive Exhibits

Whereas studies involving mock tasks generally involve a small number of users interacting over a reasonably long time (roughly from half an hour to a day), interactive exhibits of TUIs are likely to attract many users, even hundreds, many of whom will voice an opinion based on relatively short periods of interaction or even observation. This evaluation format is popular with systems that are either simple in concept or impressive in performance; systems exposed to this form of evaluation are Curlybot (Frei et al., 2000), InTouch (Brave et al., 1998), DataTiles (Rekimoto et al., 2001), and the reacTable (Jordà et al., 2007).

2.4.6 Casual Observation/No Comment

A number of papers on TUI systems include feedback based on casual observation by colleagues or visitors, often qualified with plans for future, more formal evaluation. Published accounts of Lumitouch (Chang et al., 2001), AudioPad (Patten et al., 2002), Block Jam (Newton-Dunn et al., 2003), Urp (Underkoffler and Ishii, 1999), Passive Props (Hinckley et al., 1994), Rasa (McGee and Cohen, 2001), StoryGrid (Moher et al., 2005) and Tern/Quetzel (Horn and Jacob, 2007) all fall into this category.

The other systems that were reviewed in this chapter make no comment on completed or planned evaluation.

2.5 Research Opportunities

This literature review highlights a number of opportunities for making contributions to the TUI research area. Few TUI systems have been deployed to real users in context – many claims as to the benefits of tangibility are therefore speculative, rather than grounded in real usage experiences. The majority of systems also exploit the most salient benefits of tangibility: that TUIs can mimic and augment literal domains; that tangible interaction makes actions visible to others in co-located collaboration; and that tangibility supports the creation of digital meaning in the physical world. Another gap in the TUI literature is the lack of a structured, theory-based design process. In the next chapter, I address this gap by presenting my view on how *analytic design* can be used to both explore the TUI design space and to justify how the resulting TUI design supports user activities. I apply this design process to the office context in Chapter 4, referring back to this literature review as motivation for the design and contextual evaluation of a personal desktop TUI, which supports abstract work activities through independently meaningful physical tokens.

Chapter 3

The Analytic Design of TUIs

The previous chapter demonstrates the diversity of research on interfaces and interaction styles that may be described as 'tangible'. In this chapter, I present a method for the design of TUIs that systematically moves from the design context to a contextually-appropriate design, in a manner that enables designers to bring together many different perspectives in the construction of a coherent 'story' as to how and why a design will be usable and meaningful to its future users.

3.1 The Need for an Analytic Design Process

What is often missing from presentations of TUI systems is a description of how the researchers moved from the initial problem to the final design. In many cases, insight may be cited as the origin of design, or intuition as the method of selection between competing designs. Such reasoning does not make a contribution to the body of design knowledge, however, since it fails to account explicitly for the models, theories and frameworks that are tacitly and implicitly drawn on by experienced designers. What is missing in the tangibles literature is a comprehensive yet concise formulation of the different aspects of design, along with the most appropriate conceptual and methodological tools that can be applied to the development of TUIs.

This chapter describes an analytic design process for the creation of contextually appropriate TUIs. Design is analytic when it is broken down into its elemental components, with each of these components being analysed individually. Rather than marginalizing the role of creativity in design, such an analytic approach can actively assist designers in their discovery of creative design solutions. In a manner that complements unbounded, 'blank slate' thinking, analytic design provides a vocabulary for the many and varied ways in which the different aspects of design solutions relate to both one another and to the problem context. The manifestation of creativity in such an analysis is the identification of new ways in which different aspects of a design relate to one another in ways that provide synergistic benefits for the use of that design in context. The value of analytic design, therefore, is that it provides opportunities for insight in many ways and at many levels of abstraction, as well as supplying a vocabulary for discussing design decisions that might previously have been labelled as "intuitive". The virtues of such discussion tools are described by Blackwell and Green (2003) in their presentation of Cognitive Dimensions analysis:

> *[Discussion tools] elucidate notions that are vaguely known but unformulated; they prompt a higher level of discourse; they create goals and aspirations; they encourage reuse of ideas in new contexts; they give a basis for informed critique; they provide standard examples that become common currency; and they allow the interrelationships of concepts to be appreciated.*

However, more than just a design vocabulary, I am advocating an analytic *design process*. This transition from an analytic design framework to an analytic design process is built on an identification of the channels of influence that exist between independent component analyses. Whilst any given application of the process is likely to result in a complex web of connections between analyses, there is value in attempting to tell a coherent story about how and why a design makes sense, both in its context and to its users. Even if the design process was highly non-linear, if a design cannot be rationalised as a logical progression from problem to solution then there is a fundamental problem with that design. Parnas and Clements (1985) talk about "faking a rational design process" in the context of software engineering, suggesting that designers need guidance, and will come closer to achieving a rational design if they adopt a process that encourages them to think in such a way. By analogy to mathematics, they point out that "Mathematicians diligently polish their proofs, usually presenting a proof very different from the first one they discovered [...] The simpler proofs are published because the readers are interested in the truth of the theorem, not the process of discovering it" (ibid.). Reusing this analogy here, I argue that a rational presentation of an analytically-designed TUI allows designers to best present the elegance of their solutions, even if these solutions were generated from the most inelegant of paths through the design space. Furthermore, little of general relevance would be gained by reading a detailed account of what actually happened, whist much would be lost in terms of clarity.

3.2 Overview of the Analytic Design Process for TUIs

The analytic design process for TUIs can be viewed as a rational, progressive refinement from a design context to a contextual design, i.e. one that is contextually appropriate:

1. *Context analysis* identifies the activities in a context that could benefit from TUI support – it refines a *design context* into a *design opportunity*;

2. *Activity analysis* describes the properties of a TUI that would appropriately support these activities – it refines a *design opportunity* into a *design space*;

3. *Mapping analysis* generates the physical-digital mappings of a TUI structure with these properties – it refines a *design space* into a *structural design*;

4. *Meaning analysis* provides these mappings with meaning that users can understand and adapt – it refines a *structural design* into a *surface design*;

5. *Appropriation analysis* considers the consequences of users adapting the intended meaning – it refines a *surface design* into a *contextual design*.

The linearity of this progression is for the purpose of rationally justifying design decisions, rather than stipulating how analytic design should actually proceed in practice. In the remainder of this chapter, I will present the analytic design process in detail, covering the theoretical concepts that underlie each kind of analysis, my translation of these concepts into key design vocabulary and probing design questions, and the ways in which the different analyses can influence one another. I will also discuss the potential for involving users in the process, for whilst "analytic" methods are often defined relative to "user-centred" ones, my intention here is to provide a structure for user contribution as well as designer reflection.

3.3 Context Analysis

A particularly effective way in which researchers can gain a deep understanding of a particular context is through ethnography. This involves the researcher becoming an accepted member of a social group existing in that context, participating within the group's social life over an extended period of time, and reflectively describing and interpreting the socio-cultural fabric of the context in which group activity occurs. Ethnography is both a highly time-consuming and a highly skilled practice; moreover, it is often unconstrained by the limits of specific design desiderata. In his paper entitled *Implications for Design*, Dourish (2006) argues that the "implications for design" section, concomitant with published accounts of ethnography, "may underestimate, misstate, or misconstrue the goals and mechanisms of ethnographic investigation". If this is the case, then it is reasonable to ask what we should do if our primary concerns are the contextual implications for design, with the method used to discover them being of secondary importance.

In the same paper, Dourish refers to what he claims might appear to be "discount ethnography" techniques: cultural probes (Gaver et al., 1999), technology probes (Hutchinson et al., 2003), and contextual design (Beyer and Holtzblatt, 1998). However, he argues that their focus on design represents a rejection of traditional ethnographic inquiry, rather than an abbreviation of it. Contextual design might appear to be the best compromise between in-depth ethnographic inquiry and design-oriented investigation, but it is still a relatively heavyweight method that requires many hours in the field.

3.3.1 Application of Context Analysis

What is required, therefore, is a lightweight means of capturing and presenting the characteristic features of the environment that are pertinent to the choices made by the designer.

Such a method cannot capture the rich detail of the interactions that social actors have with both artefacts in their environment and with one another, but it can give a broad-brush description of the contextual backdrop for these interactions. The approach to context analysis advocated here, therefore, is a focus on *user activities* and a high-level decomposition of context into four distinct aspects that impact upon the situated accomplishment of these activities: structural aspects, procedural aspects, cognitive aspects, and social aspects. These aspects are suitable for the early stage of design, in which investigative fieldwork can be used to uncover problems with existing support for user activities. They can be used as thematic elements of semi-structured interviews or as coding schemes for observational studies; more generally, they provide a 'quick-and-dirty' way of summarizing the ways in which context may have implications for design.

Practical application of context analysis takes the form of the following design vocabulary and design questions:

Structural Context

How are activities distributed across people, artefacts, and space?

For example, structural context in distributed software-project teams: activities are spread across many different people, places, and time-zones; problems might arise from the lack of adequate file sharing tools.

Procedural Context

How are activities initiated, co-ordinated, and completed over time?

For example, procedural context in restaurant order tracking: multiple sequential tasks are involved in the process of delivering the right food at the right time; problems might arise from the lack of adequate planning tools.

Cognitive Context

How do the cognitive demands of activities compare to the means of cognitive support?

For example, cognitive context in air traffic control: the progress of many flights must be tracked simultaneously under high levels of stress; problems might arise from the inability to set personal reminders to act at precise times.

Social Context

How do the social demands of activities compare to the means of social interaction?

For example, social context in the classroom: children of different personalities need to collaborate effectively on shared activities; problems might arise from learning resources that do support equitable participation.

3.3.2 Influence of Context Analysis

Context analysis is nominally the first stage of the analytic design process, and as such should influence all others to varying degrees. Directly, it determines the activities whose desired usability profiles will be specified by activity analysis. Indirectly, it determines the mappings that will be judged as appropriate during mapping analysis. In addition, context analysis also determines *what* things are to be represented in meaning and appropriation analysis (which respectively determine *how* these things are or could be represented).

3.4 Activity Analysis

The application of context analysis to an activity domain results in a better understanding of how appropriate different forms of interface and interaction might be for supporting the activities of that domain. The purpose of activity analysis – the next stage of the analytic design process – is to describe the abstract properties of interfaces in a manner that allows them to be compared against both one another and the requirements of the context in which they would be deployed. By viewing interfaces as notations, or abstract structures of representation and control, we can analyse the usability and suitability of those interfaces independently of their surface appearance and application semantics.

The original and best known form of such abstract analysis is the Cognitive Dimensions of Notations framework, originally created by Green (1989), and since revised and updated by Green and Petre (1996) and Blackwell and Green (2003). Cognitive Dimensions analysis has four main premises:

1. Usability is not an absolute, but a function of the activities to be performed, the notation on or through which those activities are performed, and the environment in which the notation is manipulated.
2. Usability is not a unitary scale, but a multidimensional space. Each dimension can be given a distinctive label, as is the case with the Cognitive Dimensions, with the aim of providing a shared vocabulary for design discussion.

3. Dimensions of usability trade-off against one another, so attempting to increase the usability of a notation along one set of dimensions is likely to have the side-effect of decreasing the usability of the notation along a different set of dimensions.

4. Design is the process of selecting design manoeuvres whose associated trade-offs move the notation towards the desired dimensional profile of the activities to be supported.

The core of the framework is a list of Cognitive Dimensions (CDs), which describe abstract usability properties of notations. They are at least 14 in number, although new dimensions are frequently being proposed and the set of dimensions is essentially open. The dimensions are generally neither beneficial nor harmful properties in themselves: their contribution to overall usability depends on the activities to be performed. The CDs can be easily identified in this dissertation by their typesetting convention as $\texttt{cognitive dimension}_{<\texttt{CD}>}$ – a practice we initiated in Edge and Blackwell (2006). The application of CDs has been well documented in the CDs tutorial (Green and Blackwell, 1998) and the CDs questionnaire (Blackwell and Green, 2000); however, the analytic design of TUIs is concerned not with purely digital notations, but with those that extend into the physical world.

3.4.1 Application of Activity Analysis

Rather than introducing new activities or dimensions to the CDs framework, I decided to look for particularly salient reinterpretations of the CDs that incorporated the characteristic features of physical media and the physical environment. An extensive description of the derivation of these new dimensions, which we called the *Tangible Correlates of the Cognitive Dimensions* (Edge and Blackwell, 2006), can be found in the special issue of JVLC celebrating ten years of CDs research since the paper by Green and Petre (1996). The idea was not to replace the CDs framework as an analytic tool for TUIs, but to provide a complementary perspective that focused on the physical notation. We call such notations *manipulable solid diagrams* to emphasise their tangible characteristics and abstract structural properties, at both the level of internal object configuration and the spatial arrangement of systems of objects. An extensive analysis of trade-offs in physical notations is given in the second half of our paper on the Tangible Correlates (Edge and Blackwell, 2006), in which we examine spatial and property-based options for the physical representation of relations (order and grouping) and values (discrete and categorical), as well as spatial and mechanical options for the representation and control of continuous values. Analysis of the representation options for continuous values is given as an example of Tangible Correlates analysis in Appendix A.1.

Within the main text of this dissertation, the typesetting convention of $\texttt{tangible correlate}^{[\texttt{cognitive dimension}]}_{<\texttt{TC}>}$ is used to denote the Cognitive Dimension from which the Tangible Correlate is derived. Each of our tangible correlates can be translated into a probing design question. For a TUI-based physical notation, designers should consider:

$\texttt{rootedness}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$: *To what extent do activities require low resistance to changes in the location of physical objects?*

For example, disaster response activities, such as the flood management supported by the MapTable (Reitmayr et al., 2005), may necessarily be situated in a fixed location such as a control centre. However, activities such as the media management supported by MediaBlocks (Ullmer et al., 1998) may require greater physical portability.

$\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$: *To what extent do activities require the preservation of physical structures for later inspection?*

For example, tangible formulations of database queries, such as Ullmer's (2002) tangible query interface, might be kept for later use or reuse, but in turn such representations may need lower $\texttt{rootedness}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$ in order to be transported to archival locations.

$\texttt{shakiness}^{[\texttt{error proneness}]}_{<\texttt{TC}>}$: *To what extent do activities require protection against difficult to reverse changes of physical state?*

For example, physically-creative activities such as the landscape sculpting of Illuminating Clay (Piper et al., 2002) may require a safeguarding of the physical creation, but storage problems can occur if such creations also require high levels of $\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$.

$\texttt{purposeful affordances}^{[\texttt{role expressiveness}]}_{<\texttt{TC}>}$: *To what extent do activities require the visible action possibilities of objects to be computationally interpreted?*

For example, activities based on tangible communication, such as the Lumen shape displays (Poupyrev et al., 2004), may require the majority of visible action possibilities to be interpreted symbolically, but this could also increase $\texttt{shakiness}^{[\texttt{error proneness}]}_{<\texttt{TC}>}$ if these actions are easy to perform accidentally.

$\texttt{bulkiness}^{[\texttt{diffuseness}]}_{<\texttt{TC}>}$: *To what extent do activities require spatial representations that extend in three dimensions?*

Tangible elements for media access, such as the MediaBlocks (Ullmer et al., 1998), do not need to extend far in three dimensions because physical composition is not required for meaning creation. However, the resulting interface may not exhibit the same level of $\texttt{purposeful affordances}^{[\texttt{role expressiveness}]}_{<\texttt{TC}>}$ as interfaces in which tangibles are visibly designed to work together.

$\texttt{structural correspondence}^{[\texttt{closeness of mapping}]}_{<\texttt{TC}>}$: *To what extent do activities require physical and digital information structures that resemble one another?*

For example, a tangible interface used to visualise the forces within physical structures, such as Senspectra (LeClerc et al., 2007), may benefit from coincident physical and digital representations, but such structures also exhibit more $\texttt{bulkiness}^{[\texttt{diffuseness}]}_{<\texttt{TC}>}$ than interfaces based on independently meaningful elements.

$\texttt{juxtamodality}^{\texttt{[juxtaposability]}}_{\texttt{<TC>}}$: *To what extent do activities require decoupling between observation and manipulation?*

For example, activities in which localised actions have more globally-observable effects may require eyes-free operation, such as the physical tools used to adjust the simulated time of day and wind direction in the Urp interface for urban planning (Underkoffler and Ishii, 1999). However, such approaches result in lower $\texttt{structural correspondence}^{\texttt{[closeness of mapping]}}_{\texttt{<TC>}}$ since physical and digital representations are not coincident.

$\texttt{hidden augmentation}^{\texttt{[hidden dependencies]}}_{\texttt{<TC>}}$: *To what extent do activities require digital augmentation that is physically obvious?*

For example, a clearly augmented TUI may improve its legibility to novices, such as Carvey et al.'s (2006) Rubber Shark as User Interface based on weight indexing with electronic scales, but such legibility is difficult to achieve in interfaces that use a high degree of $\texttt{juxtamodality}^{\texttt{[juxtaposition]}}_{\texttt{<TC>}}$.

$\texttt{rigidity}^{\texttt{[viscosity]}}_{\texttt{<TC>}}$: *To what extent do activities require low resistance to changes in the configuration of physical objects?*

For example, speed-critical activities such as musical improvisation on the reacTable (Jordà et al., 2007) may require low resistance to physical reconfiguration, but this may require an increase in $\texttt{hidden augmentation}^{\texttt{[hidden dependencies]}}_{\texttt{<TC>}}$ relative to interfaces whose workings are physically explicit (such as using physical connection rather than perceptual line-up).

$\texttt{unwieldy operations}^{\texttt{[hard mental operations]}}_{\texttt{<TC>}}$: *To what extent do activities require low levels of physical manipulation difficulty?*

For example, physical manipulation difficulty may impede certain activities, such as simultaneously and independently manipulating many articulated joints of a Topobo construction (Raffle et al., 2004), but the decomposition of complex manipulative functions into simpler component parts may also increase the degree of $\texttt{rigidity}^{\texttt{[viscosity]}}_{\texttt{<TC>}}$ for expert users.

As well as considering the extent to which activities require their supporting notations to exhibit each of these properties, designers should also consider to extent to which the context would support notations exhibiting those properties judged as desirable. For instance, whilst a tangible construction with high $\texttt{structural correspondence}^{\texttt{[closeness of mapping]}}_{\texttt{<TC>}}$ may generically support the local creation and sharing of information, certain contexts may preclude the high degrees of $\texttt{bulkiness}^{\texttt{[diffuseness]}}_{\texttt{<TC>}}$ that complex physical structures can demonstrate.

These questions should not be confused with a checklist: they are a means of making the underlying Tangible Correlates understandable. As with the Cognitive Dimensions, changes to a TUI design are likely to be accompanied by a corresponding set of conceptual trade-offs between these abstract usability properties. At this stage of the design process however, they are simply used to describe the nature of activities to be performed.

3.4.2 Influence of Activity Analysis

The activity analysis presented in the previous section helps to characterise an abstract design space in terms of desirable usability properties. Only part of this space will be considered usable relative to the activities to be performed on or through the notation, and the environment in which that notation is manipulated. The next logical step in the analytic design process – mapping analysis – explores the usable region of this space through concrete structural designs which can then be compared to this desired dimensional profile. When tangible prototyping is adopted as a means of physically exploring this design space, prior performance of activity analysis will help ensure that it is as an elaboration of design ideas that are fundamentally sound with respect to the desired usability profile. An appreciation of the trade-offs associated with different physical instantiations of common information structures, such as those we presented in Edge and Blackwell (2006), means that unsuitable designs can be ruled out from the start.

It is also possible for activity analysis to reveal gaps in understanding about the context – for instance, whilst a particular activity might benefit from representations with high $\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$, it might not be immediately apparent to what extent the context would permit the persistent display or archiving of physical structures.

3.5 Mapping Analysis

Whilst digital notations are highly unconstrained and highly conventionalised, physical notations are the exact opposite. Physical objects do not have the same indeterminate and malleable potential of visual digital imagery – they are constrained by gravity, materiality and mechanics. The physical design space is sufficiently constrained that *descriptions* of physical notations can be analysed. These descriptions naturally take the form of those structures and actions that are interpreted as meaningful by the system – i.e. its syntax. For TUIs, many possible syntactic structures are drawn from the "variables" of physical objects, building on previous work on graphical variables by Bertin (1967), MacEachren (1995), and Engelhardt (2002). These "physical variables" (see Table 3.1) form the building blocks of TUI syntax, as well as providing opportunities to create $\texttt{secondary notation}_{<\texttt{CD}>}$ – extra information in means other than formal syntax.

What distinguishes the physical variables of TUIs with those of static diagrams, animated diagrams, and haptic user interfaces is the ways in which these variables can be changed over time. Whereas static diagrams exhibit no temporal properties other than persistence (and perhaps changes in the location and orientation of the diagram itself), animated diagrams can adjust all of these visual variables over time. Furthermore, whilst haptic user interfaces can dynamically and automatically adjust many of these tactile and mechanical variables as well as some visual variables, tangible user interfaces rely more on user control. This is not to say that TUIs cannot contain haptic or visual-temporal elements, however – the PICO system (Patten and Ishii, 2007) is a good example of a TUI that can automatically

Physical Variable Type	Variables
Spatial Configuration	Location (relative to other objects or within a space); Orientation (relative to other objects or within a space);
Visual Attribute	Size (height, width, depth, volume); Shape (structure, symmetry, curvature, etc.); Colour (hue, saturation, brightness); Clarity (crispness, resolution, transparency); Augmentation (annotation, adornment)
Tactile Attribute	Material (smoothness, temperature); Texture (directionality, size, density); Mass (gravitational, angular); Deformability (elasticity, plasticity, viscosity)
Mechanical Configuration	Resistance (compression, tension, shear, bending, torsion)

Table 3.1: Physical Variables used in TUI Syntax

adjust the location of its tangibles, whilst the digital representations of many TUIs can be seen as 'dynamic diagrams'.

Beyond considerations of purely physical variables, however, mapping analysis is based on comprehensive description of how selected physical variables create a physical syntax that maps onto a digital information structure within a TUI. It can be used to both explain the general character of a particular TUI, and to explore a design space by systematically considering manipulations of the various aspects of mapping. There are four such aspects – spatial, action, attribute, and temporal – that combine to cover the nature of the physical-digital relationship as it extends over space and time, and representation and control.

3.5.1 Application of Mapping Analysis

Spatial Mapping

How are the physical configurations of objects computationally interpreted?

Spatial mapping refers to the relationship between the physical arrangement of objects and their digital interpretation. There are three conventional "styles of mapping" in TUIs (Ullmer and Ishii, 2001); these are implicitly spatial mappings that can be referred to as *'pure' spatial*, *relational*, and *constructive*.

Firstly, in a 'pure' spatial mapping, the physical objects model a target domain in which their spatial arrangement is interpreted literally. A prototypical example of a pure spatial interface is a map, which represents some territory somewhere in the world. The MapTable

(Reitmayr et al., 2005) is an example of a TUI that augments paper maps through projected digital overlays and physical points of command and control.

Secondly, in a relational mapping, object positions are not interpreted relative to some absolute frame of reference (as is the case with pure spatial interfaces), but relative to one another. This added flexibility allows more abstract relationships to be expressed spatially, such as order and grouping. The Senseboard (Jacob et al., 2002) is an example of a TUI that represents these relational concepts within the constraints of a two-dimensional grid, for the grouping of related papers into conference sessions, and ordering within and between those sessions in order to create a conference schedule.

Finally, in a constructive mapping, object positions are also interpreted relative to one another. Rather than expressing abstract relations between objects, however, the composition of objects is interpreted as a higher-level structure. Like pure-spatial mapping, constructive mapping is usually applied to reality-oriented domains (as opposed to abstraction-oriented domains); like relational mapping, constructive mapping has an internal frame of reference. The Topobo system (Raffle et al., 2004) is an example of a TUI that uses passive construction elements to build biomorphic structures, and active construction elements to give them kinetic memory, allowing the manipulations of those elements to be reproduced in a "playback" mode. In constructive mappings, meaning is derived from the whole, rather than from the component parts.

Action Mapping

How do physical actions lead to digital effects, in terms of timing, location, and similarity?

Action mapping refers to the relationship between physical actions and digital effects. There are two main characteristics of this relationship, which I will refer to as *indirection* and *compatibility*. These are both borrowed from the post-WIMP *Instrumental Interaction* work of Beaudouin-Lafon (2000), with the term "interaction instrument" describing "the association of a physical part (the input device) and a logical part (the digital representation on the screen)". Although this interaction paradigm runs somewhat parallel to tangible interaction, the concepts remain readily applicable.

As part of a TUI's action mapping, indirection denotes the spatiotemporal offsets between physical actions and digital effects. For example, the AlgoBlock system (Suzuki and Kato, 1993) exhibits high indirection because the effects of physical construction are seen as delayed program execution on a decoupled display screen, whereas the IP network simulator (Kobayashi et al., 2003) running on the Sensetable (Patten et al., 2001) exhibits low indirection, since real-time digital representations are projected around the physical control pucks. The concept of indirection is mirrored in Fishkin's taxonomy (Fishkin, 2004) by the notion of "embodiment", which is categorised into decreasing levels of embodiment: full, nearby, environmental, and distant.

The second characteristic of action mapping, compatibility, denotes the match between physical actions and digital effects. In my interpretation of compatibility for TUIs, I make a further differentiation between the compatibility of media, and the compatibility of marks. Media compatibility results when actions and effects are perceived in the same way, either through vision, touch, or sound. Mark compatibility, on the other hand, refers to the structural similarity between actions and their effects – this is called "isomorph effects" by Hornecker and Buur (2006). The choice of action mappings in general will depend on the goals and structures of the activities to be performed, in their context of performance.

Attribute Mapping

How do physical object properties relate to digital information attributes, in terms of the goodness, multiplicity, and persistence of the association?

Attribute mapping refers to the relationship between the physical properties of objects and the attributes of the digital information they represent. There are two main characteristics of this relationship, which I will refer to as *integration* and *coupling*.

Integration is the third of the three properties of "interaction instruments" as laid out by Beaudouin-Lafon (2000), along with indirection and integration. It refers to the match between the physical degrees of freedom of the input devices, and the dimensions of the digital attribute values they represent and control.

A closely related concept is coupling, with which I refer to the permanence and multiplicity of the associations between the physical and the digital. These concepts were used by Fitzmaurice (1996) in his dissertation on Graspable User Interfaces, in which he represented coupling as a continuum from "time-multiplexed" to "space-multiplexed". In a purely time-multiplexed interface, for example a mouse-based WIMP interface, a single physical device (the mouse) transiently acquires, manipulates and releases virtual functions (GUI windows, icons, menus, and other widgets) in strict time sequence, because there is only a single point of control. In a purely space-multiplexed interface, for example a keyboard-based command interface, multiple physical devices (keyboard keys) are permanently coupled to virtual functions (symbolic characters), offering multiple points of control. The concept of Graspable User Interfaces builds on this notion of space-multiplexing, and adds the additional characteristics of device specialisation, concurrency, spatial awareness, and spatial reconfigurability (ibid.). Whilst the majority of TUIs based on systems of objects implicitly aim to achieve extreme spacemultiplexing, it remains instructive to consider the opportunities created by relaxing the constraints on integration and coupling.

A concept that conflates all of the characteristics of action and attribute mapping – integration, coupling, indirection, and compatibility – is "degree of coherence" from Koleva et al. (2003), which refers to the degree to which "linked physical and digital objects might be perceived as being the same thing". Whilst a useful concept, it is too high-level for the detailed analysis of mapping considered here.

Temporal Mapping

How does physical specification of behaviour at one time lead to digital behaviour at a later time?

The final aspect of mapping is temporal mapping, which refers to the relationship between the physical specification of behaviour and its digital execution. Most metaphors of interaction implicitly build on Shneiderman's (1987) concept of "direct manipulation", in which the current state of the interface gives complete feedback about users' progress towards their goals, and in which the effects of actions are immediately visible. However, it is not always the case that users need to construct concrete representations that describe the present state of the world; in many instances, users need to construct abstract representations that specify some future behaviour. The current state of the world is not all that matters – the ultimate effects of users' actions are contingent on both the potentially complex interactions between future events, and the 'activation' of the representational rules or procedures constructed by the user. In software, this interaction metaphor would be recognised as scripting or programming, but in the physical world, such specifications of behaviour are manifest in *manipulable solid diagrams* of the kind discussed in Section 3.4.1. Under this metaphor, the specification of future behaviour can derive either from the dynamic manipulation of solid elements (hence "manipulable"), or from their static structure (hence "diagram").

Perhaps the best known example of *dynamic* temporal mapping in TUIs is in the Topobo system (Raffle et al., 2004), which utilises programming by demonstration: the manipulations performed by a user in record mode are temporally translated to become the behaviour of the structure in play mode. An example of *static* temporal mapping in TUIs is the MediaCubes system (Blackwell and Hague, 2001), in which networks of domestic appliances can be programmed through the capture of physical cube arrangements. Whilst much TUI design focuses on direct manipulation, the use of tangibles as components of manipulable solid diagrams is a relatively unexplored and potentially fruitful area of research.

3.5.2 Influence of Mapping Analysis

Mapping analysis is most closely related with the activity analysis that precedes it, in that the descriptions of TUI mappings are evaluated against the profile of Tangible Correlates (and perhaps also the profile of Cognitive Dimensions) that was generated during activity analysis, and that *design manoeuvres* from one set of mappings to another should be informed by the associated trade-offs in terms of usability dimensions. The process could also be done in reverse, however, with tangible prototyping used as a participatory design technique to elicit from potential the suitability of various interface structures for the contextual activities that could be performed. Whilst this approach attempts to capture context as efficiently as possible through joint design investigation with real users, it relies on participants being able to articulate their needs and to see beyond the capabilities of familiar technologies – something that is not always the case. Particularly with emerging technologies such as TUIs, part of the designer's responsibility is to create compelling interactive prototypes capable of convincing potential users of the potential benefits of such technologies.

As well as referring back to previous logical stages of the analytic design process, mapping analysis also provides a structure on which the next stage, meaning analysis, can be applied. Another potential feed-forward influence of mapping analysis is to the final logical stage of appropriation analysis. Those physical variables that are not used in the formal syntax of the system are available for use as `secondary notation`$_{<\mathrm{CD}>}$ – it is these physical variables that can be appropriated by users to represent extra information that is not expressible within the formal syntax of the system.

3.6 Meaning Analysis

The mapping analysis of the previous section provides a set of conceptual tools for the examination and description of the structure of a concrete TUI design. However, the analysis is still abstract in that it deals with interactional structures rather than representational forms. In particular, mapping analysis does not deal with the question of how users find the TUI to be meaningful.

This issue is addressed by performing meaning analysis, the penultimate stage of the analytic design process. This looks at the relationship between what is perceived by the user and what was meant by the designer – it is about the legibility of a design. By this point, previous analytic stages should have validated all but the surface design of the TUI. Meaning analysis helps the designer tell a plausible story about how users will interpret the perceptual and conceptual rendering of the interface as a meaningful system.

3.6.1 Application of Meaning Analysis

Visual meaning

How does the visual appearance of each interface object suggest action possibilities?

Visual meaning refers to the relationship between the visual appearance of an object or set of objects, and the perceived action possibilities. It is closely related to the concept of *affordance*, coined by Gibson (1979) in his work on *The Ecological Approach to Visual Perception*: "The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or for ill". From the perspective of design, Gibsonian affordances are those interactions that are possible with an object, in its context of use, by a human user. The concept of affordance was picked up by Norman (1988), subtly changed, and introduced to the field of HCI in his book *The Psychology of Everyday Things*, in which he says that the affordances of objects "convey messages about their possible uses, actions and functions". However, this refers to the perceived affordances of objects, which may differ from their actual affordances (which may be illusory or go unnoticed). Bærentsen and Trettvik (2002) tie the two concepts together by identifying the fact that artefacts have both possible and intended uses, saying that the task of design is to draw attention to the uses that are intended in the context of applicability of the artefact. It is this matching of user perception and designer intent that visual meaning analysis aims to address.

It is important to realise that affordances are not 'magical' properties of objects, but cues for the 'animal' that relate current object perceptions to previous interaction experiences.

We turn knobs not only because they appear unyielding to other actions such as pushing, pulling and bending (even if this is not the case), but because we have learned through experience that the most common way to achieve the desired effect is through turning. If turning fails, then we may become aware that our expectations and learned behavioural patterns have been violated – in such situations, the actual affordances of the object do not match their perceived affordances. When 'designing affordances', it is therefore essential to take into account the history of users' interactions in the world.

In the world of TUIs, many complementary affordances are supplied by the structural form of the interface, conventionally belonging to one of three categories: interactive surface, constructive assembly, and token+constraint (Ullmer and Ishii, 2000). In *interactive surface* systems, the primary affordances are the placement, orientation, and movement of physical objects on an underlying surface, which provides dynamic digital representations that change based on object configuration. In *constructive assembly* systems, the primary physical affordances are based on the means of connection between objects, from visible material (e.g. Velcro) and mechanical (e.g. plug and socket) connectors, to connection based on 'invisible' physical phenomena (e.g. magnetism and suction). However, it is in *token+constraint* systems that affordances are most utilised: "Even when tokens and constraints are physically separated, their physical complementarity to each other enable them to passively express allowable combinations and alternative usage scenarios" (Ullmer et al., 2005).

Haptic Meaning

How does the haptic experience of using each object guide action performance?

Haptic meaning refers to the relationship between haptic experience and the perception and interpretation of intended action performance. This is closely related to Norman's (1988) description of *feedback*. For example, a steering wheel looks like it could be used for turning with both hands, but its appearance alone conveys little information about its qualities of use. The experience of manipulating a steering wheel gives the user further cues as to how it should be used; for instance, the weightiness of a car's steering mechanism and its self-centring characteristics suggest that movements should generally be slow and deliberate. Visually suggesting action possibilities to the user is only half of the solution – guiding their action performance in the appropriate manner may require many iterations of physical design and test. It is likely that the modelling material of the Illuminating Clay system (Piper et al., 2002) underwent a number of design cycles to find the right balance between ease of intentional sculpting and difficulty of accidental damage.

As with visual meaning, it is through a token+constraint structural approach that TUIs best leverage haptic meaning. Ullmer et al. (2005) list the following benefits of the token+constraint approach directly associated with enhanced haptic feedback:

1. *increased passive haptic feedback;*
2. *increased prospects for active force feedback;*

3. *decreased demands for visual attention;*
4. *increased kinesthetic awareness.*

However, it is not just in token+constraint systems that haptic feedback is important. In any TUI, the size, shape, weight and surface properties of tangibles are all important channels for the communication of intended action performance to the user. For example, the action of sliding a tangible across a surface is strongly affected by the weight of the object and the frictional properties of its surface, whilst the re-orientation of a hollow cube in three dimensions will have a different 'feel' depending on the volume and viscosity of any fluid contained within it.

Functional Meaning

How does the physical form of each object signify its conceptual function?

Functional meaning refers to the relationship between physical form and conceptual function, or how users are guided to interpret the semantics of atomic tangible elements in a particular way. This falls within the field of *semiotics*, or the science of signs. A sign can be defined as "something which stands to somebody for something in some respect or capacity" (Cobley and Jansz, 2004). This definition is based on the triadic sign relations of Charles Sanders Peirce (1958), in which something – the *representamen* – stands for something else – the *referent* – in a relation that entails a certain 'result' – the *interpretant* – in the mind of the observer.

Peirce argues that just as there are three formal aspects of a sign (representamen, object, and interpretant), there are also three fundamental categories of phenomena: *qualities*, *brute facts*, and *laws*. These result in three different ways in which the physical form of a tangible element (representamen) can stand for its conceptual role within the interface (referent). Considered from the perspective of TUI design, two of these – qualities and laws – relate to signification by atomic tangible elements (the third, "brute facts", relates to signification by relations between tangible elements and will be discussed in the next section).

In *iconic* signification, the representamen stands for its referent through similarity of *qualities*, where similarity may be literal, analogical, or metaphoric. The tangible archetype of a literal icon is a physical model that mimics an element of the conceptual domain; an example of such *literal tangible iconicity* is a building from the Urp system for urban planning (Underkoffler and Ishii, 1999). The tangible archetype of an analogical icon is an articulated solid diagram, whose abstract structure mimics the structure of an element of the conceptual domain, having analogous relations of parts; an example of such *analogical tangible iconicity* is the nodes and links of the Senspectra system (LeClerc et al., 2007). Metaphoric icons have no such tangible archetype, but draw on shared functional characteristics or structural analogies from other conceptual domains.

In *symbolic* signification, the representamen stands for its referent through arbitrary or conventional rules or *laws*. Such symbolic relationships need to be learned in order to be meaningfully interpreted. Whilst symbols are generally less intrinsically expressive than icons and indices, strong cultural symbols are often much more compact than their iconic counterparts.

Relational Meaning

How does the spatial configuration of objects signify conceptual relations?

Whereas functional meaning refers to the signification of conceptual functions by atomic tangible elements, relational meaning refers to the signification of conceptual relations by the spatial arrangement of tangibles. This is not necessarily a case of what spatial relations the TUI can detect, but reasoning how and why users will interpret these relations as meaningful.

The way in which humans perform abstract reasoning about the world they inhabit is the focus of Lakoff's (1987) *spatialization of form hypothesis*, which states that we structure our bodily experiences preconceptually using kinaesthetic image schemas, and that we map these image schemas metaphorically to the conceptual structures we use in abstract reasoning. Such structures – including *container* schema, *part-whole* schema, *link* schema, *centre-periphery* schema, *source-path-goal* schema, *front-back* schema, *up-down* schema and *linear order* schema – form the basis of Lakoff's cognitive semantics and provide a theoretical basis for the design of syntactic relations in TUIs.

A corollary for tangible user interfaces, by virtue of the structural isomorphism underlying the notion of metaphor, is that the same metaphors can be used to map these abstract conceptual structures back into image schemas, which we can then project onto our arrangement and configuration of objects in the world. Thus image schemas, the means by which we structure our bodily experiences and render them meaningful, are also the means by which we create meaning in the world. Whether through the arrangement of paper, books, sticky notes or a variety of other media, including the elements of tangible interfaces within or without the sensing apparatus, it is the use of image schemas through which we structure knowledge in our heads and in the world around us. The potential use of image schemas in TUI design had also been advocated by Hurtienne and Israel (2007), in which they derive "metaphorical extensions" from traditional image schemas that are directly applicable to tangible structures. The essential point they convey is that designers should be able to explain their constructions in terms of image schemas and their metaphoric projection onto TUI syntax. In this dissertation, I am proposing a broader consideration of meaning in TUIs that is not confined to the formal syntax. In the words of Dourish (2001), "it is important not to imagine that the application's boundaries contain everything that matters".

One particular way in which users can create meaning in the world with the elements of TUIs is to place them on, adjacent to, or around meaningful objects or places in the local physical environment. Such *indexical* signification is the third and final way in which a representamen can stand for its referent, and does so through the *brute facts* of causation (Cobley and Jansz, 2004), which direct attention to a particular spatiotemporal region. Classic examples of indices are clock faces and thermometers, which demonstrate the influence of time and temperature respectively. Another classic example is a finger pointing towards the moon: the representamen (the pointing finger) stands for a region of the sky (the referent), from which the concept of the moon is triggered (the interpretant). Placing a tangible element in a particular relation to an existing physical object or place makes a 'pointing finger' out of the tangible: the tangible *in relation to* the object or place forms the representamen, which stands for the conceptual association between the two, which in turn entails cognizance of the placement act that brought the two into a state of mutual signification.

3.6.2 Influence of Meaning Analysis

The ways in which tangible objects and their relations are rendered meaningful within the confines of the system also largely determines the ways in which they may be given additional or alternative meanings. It is this ability and propensity of users to create their own meaning in the environment that forms the basis of the final logical stage of the analytic design process: appropriation analysis.

The design decisions of how to represent object functions and relations in meaningful ways rely on a good understanding of what systems of meaning already exist for potential users in context. The process of meaning analysis may uncover gaps in this understanding that can only be filled through additional context analysis.

3.7 Appropriation Analysis

Appropriation analysis considers how users will adopt and adapt the TUI to suit their individual and collective needs, determined by their interaction context. It recognizes that whilst the nature of this context should help guide the process of design, the nature of TUI deployment is such that embedded TUI use will have a reciprocal effect that can subtly yet significantly transform the interaction context. The two primary motivations for appropriation are to ease the cognitive burden of activities and to mediate social interactions.

It is instructive to think of the appropriation of tangible interface elements as the creation of cognitive and social information structures out of the opportunities for `secondary notation`$_{<\text{CD}>}$ provided by the formal TUI notation. The goal of the designer is to capture only as much of the notation space as is necessary to provide a useful system. Capturing and interpreting the marks made by users on the surfaces of tangible objects, as well as the more abstract 'marks' made by their actions in space and over time, reduces the amount of `secondary notation`$_{<\text{CD}>}$ provided by the notation, and hence the degree of appropriation possible. The extent to which the formal syntax of a notation should be traded-off against the remaining `secondary notation`$_{<\text{CD}>}$ is an issue of the *appropriateness* of such potential appropriations in the intended context of use.

3.7.1 Application of Appropriation Analysis

Cognitive Appropriation

How can users adapt the interface to create external cognitive structures?

The cognitive appropriation of TUI elements is essentially a means of *external cognition* (Scaife and Rogers, 1996), the theory of which is based on the idea that cognition encompasses both internal representations "in the head", and external representations "in the world". There are three main cognitive benefits to be derived from using external representations and tools as resources for general cognitive activities (Preece et al., 2002):

1. *Externalising to reduce memory load* can offload the burden on memory over time by reminding people what to do, or when to do it – e.g. scheduling timed alerts on a PDA.
2. *Computational offloading* can simplify a problem by using more appropriate external representations and tools – e.g. from mental arithmetic, to pen-and-paper arithmetic, to calculator-based arithmetic.
3. *Annotating and cognitive tracing* respectively refer to the modification of external representations or the spatial relationships between them – e.g. managing the items in a paper-based list or rearranging a hand of cards.

These ways of appropriating the physical environment for cognitive ends are the target of much research in Cognitive Science; in particular, they have been addressed in detail by the work of David Kirsh and Don Norman. Externalising to reduce memory load is a form of *jigging* or *informationally structuring* the environment (Kirsh, 1995), used to provide *entry points* for future activities (Kirsh, 2001). Computational offloading can be supported by *epistemic actions*, which are "physical actions that make mental computation easier, faster, or more reliable" (Kirsh and Maglio, 1994). Cognitive tracing makes use of *spatial organisation strategies*, classified by their goal of aiming to "simplify choice", "simplify perception", or "simplify internal cognition" (ibid.), and providing an *activity landscape* – "the construct resulting from users project structure onto the world, creating structure by their actions, and evaluating outcomes" (Kirsh, 2001). All three types of benefit are the result of viewing physical objects as *cognitive artefacts*, or external representations of "knowledge in the world" (Norman, 1991); such artefacts can provide *memory aids* to reduce memory load (Norman, 1988) as well as *transform the task* to simplify the problem (Zhang and Norman, 1995).

In the analytic design process, these three types of cognitive benefit are transformed into questions that probe potential cognitive appropriations:

1. How does the TUI support *cognitive externalisation* – in what ways might tangibles be used as prospective memory aids, or reminders?

2. How does the TUI support *cognitive simplification* – in what ways might tangibles be used to solve problems in the world, rather than the head?
3. How does the TUI support *cognitive tracing* – in what ways might tangibles be annotated or marshalled into and between meaningful configurations?

Considering such questions before the design of a system is finalised can ensure that its designed affordances are tailored towards the kind of cognitive appropriations that should be encouraged, as well as to ensure that the design does not unnecessarily preclude such appropriations, for example by using materials that cannot be annotated, or cannot be moved outside of the sensing region of the TUI.

Social Appropriation

How can users adapt the interface to create external social structures?

The social appropriation of TUIs and their elemental parts is related to their symbolic value and the existing systems of meaning in the given context. The sociological theory of *symbolic interactionism* (Blumer, 1969) is used here in the same manner as external cognition was used above, to create questions probing potential social appropriation of tangible interface elements; just like external cognition, it too has three main premises:

1. *Human beings act toward things on the basis of the meanings that the things have for them.*
2. *The meaning of such things is derived from, or arises out of, the social interaction that one has with one's fellows.*
3. *These meanings are handled in, and modified through, an interpretive process used by the person in dealing with the things he encounters.*

These are translated into probing questions as follows:

1. How does the TUI affect *social actions towards things* – in what ways might people act towards tangibles based on their functional but also their symbolic social meaning?
2. How does the TUI affect *social production of meaning* – in what ways might tangibles be used to create new kinds of meaning in ways that didn't exist or weren't apparent before?
3. How does the TUI affect *social interpretation of things* – in what ways might individuals change their interpretation of tangibles over changes in place, time, and situation?

Answering these questions will result in a better understanding of how actual usage of the system might differ from that originally intended. As with cognitive appropriation, it is up to the designer to determine the degree to which potential social appropriations are communicated to users. However, it is important to realise that whereas external cognition is *a way* of doing *certain things* in a manner that produces cognitive benefits, symbolic interactionism is a theory describing *the way* we do *everything* as social creatures. Whereas potential cognitive appropriations are relatively straightforward and limited, social appropriations are more subtle and varied. What is important, therefore, is to ensure that the TUI design takes into account any social appropriations that can be anticipated in advance, even if they may only develop organically over an extended period of time.

3.7.2 Influence of Appropriation Analysis

Appropriation analysis, although nominally the final stage of the analytic design process, can develop insights that can affect any of the preceding stages. It might become apparent that the systems of meaning embodied by the interface could lead to undesirable interpretations of certain actions, triggering a re-visitation meaning analysis. For example, choosing to physically represent members of social networks such that users can organise and display their contacts in their desktop environment could lead to competitive behaviour in terms of who has the most 'friends', detracting from the potential goal of giving more direct access to communication with the select few who have an immediate bearing on current work activities.

Useful appropriations might also be identified, but precluded by the choice of mappings – leading to a similar re-visitation of mapping analysis. Appropriation analysis could also identify non-syntactic uses of the interface that would benefit from computational support, in which case these new activities would need to be incorporated into the activity analysis. Finally, designers may realize that they don't fully understand the implications of certain appropriations, resulting in a need for further context analysis. This kind of iteration is to be expected in the design of any TUI for real use. Rather than analysis replacing user-centred techniques, it should instead help to make the issues clearer, such that more insight can be gained from each session of user involvement.

Chapter 4

Designing for Peripheral Interaction

My choice of application domain – the desk-based, office context – was primarily influenced by two factors. Firstly, a review of previous TUI research (see Chapter 2) revealed a strong focus on shared interfaces in which users jointly engaged in co-located, collaborative activities; however, little work had been done on supporting the generic office worker at their desk. Secondly, existing ethnographic studies on the use of paper in work environments (to be discussed in Section 4.1.3) have shown it to play an important role in the coordination of social activity. I therefore became interested in how TUIs could support both individuals and groups in the office context, by exploiting the advantages of both physical and digital media.

In order to influence the design of a TUI that could exploit opportunities for office-based tangible interaction, I conducted a number of interviews and observation sessions at the offices of a large, multinational technology company. These interviews were the basis of the context analysis presented next – the aims and themes of which were presented in Section 3.3. The remainder of this chapter is a presentation of the analytic design process as I applied it to the office context.

The work required to conduct this investigation falls within the field of HCI and is typically interdisciplinary, involving aspects of computer science, engineering, psychology, sociology, and design. A focus on users is maintained throughout this dissertation; firstly through investigative fieldwork into how existing office work could be improved through the introduction of augmented physical tokens, and secondly through the evaluation of the resulting design via an extended deployment in an organizational setting.

4.1 Context Analysis

To investigate potential design opportunities for the desk-based, office environment, I visited a multinational technology company structured around small, co-located project teams, and interviewed representatives of different managerial strata: an engineer, a project manager, and a senior manager in charge of coordinating project managers and their teams. The interviews uncovered a number of problems with existing work processes and practices. One of the more subtle difficulties was that the use of email had pervaded almost all areas of working life, addressing many functions that had, in the past, been dealt with by face-to-face conversation.

The benefits of email, in terms of providing a record of communication for future reference and accountability, as well as lowering communication overheads through its direct and asynchronous nature, meant that email had become the de facto standard for workplace communications even when its appropriateness was questionable. This had resulted in a state of email overload, in which information often became hidden, lost, or forgotten due to inbox messages scrolling off the screen, compounded by the homogeneous appearance of emails, and complicated by the personal systems of filing and labelling used to organise them.

A second problem associated with the default use of email for office communication was that people didn't talk to one another as much as they used to – not just about particular issues, but about work in general. Combined with the standard practice of weekly project meetings, this had resulted in a general lack of awareness about the work status of other team members. These were formally updated once a week in time for the project meeting, but rarely at other times, and to compound matters further, status updates that did occur were collected by email.

Other problems were related to the retrospective and inaccurate completion of timesheets; the inability to easily share information recorded on physical media such as note books, whiteboards, and post-its; and the inappropriateness of planning work in calendars, which failed to reflect the reality of how work was carried out: schedule, start, suspend, restart, reschedule, and so on. In all cases, the problems appeared to stem from the interactional and attentional costs of creating and updating digital information structures about work, in parallel with actually doing it. Such auxiliary work activities, of which email management, timesheet completion, information sharing, and work planning are all instances, can often become marginalized, neglected, or forgotten due to the pressures of multitasking.

4.1.1 Motivations from Literature Review

TUIs are still a nascent research area, and most research to date has exploited the most salient advantages of physicality. Following on from early proofs of concept such as the metaDESK (Ullmer and Ishii, 1997) and mediaBlocks (Ullmer et al., 1998), recognition that interaction with tangibles makes actions visible to others, who can interact concurrently with the multiple points of control that TUIs make explicit, has led to a strong research emphasis on tabletop TUIs for co-located collaboration. Application domains have been hand-picked for their collaborative requirements, resulting in a broad range of tabletop TUIs for urban design (Underkoffler and Ishii, 1999), music performance (Jordà et al., 2007), storytelling (Gorbet et al., 1998), and video editing (Zigelbaum et al., 2007), amongst others.

However diverse these projects appear, they share a number of features that characterise the current state of play in TUI research. In the "shareables" work of Hornecker et al. (2007)

and beyond, users generally perform a *focal activity* around a *single, shared interface*, which employs a *spatial syntax* to interpret the relative configuration of tangibles. Each of these features is predicated on the superiority of tangibles for the representation and control of information structures belonging to inherently spatial domains. These domains are chosen because of their suitability for a spatial syntax – they meet the fundamental criterion set out in the influential *Emerging Frameworks* paper that "the physical state of the interface artefacts partially embodies the digital state of the system" (Ullmer and Ishii, 2001). Such rigid spatial interpretation precludes having more than one such interface, since in order to maintain representational consistency, changes made to any one interface would need to be replicated in all of the others. Whilst the actuated workbench project (Pangaro et al., 2002) and subsequent incarnations demonstrate the possibility of spatial actuation in TUIs, whether or not it is desirable or even feasible on a large scale remains an open question. As such, the existence of a spatial syntax generally leads to a single, shared interface, and the lack of a motivating spatial syntax – as is the case with the abstract objects of interest in the office environment – has generally resulted in a lack of research into appropriate TUI support.

4.1.2 Insight from Abstract Analysis

The conventional styles of mapping for TUIs, whether pure-spatial, relational, or constructive, all impose a *spatial syntax* on the placement of physical elements, defining how any arrangement will be interpreted by the system. This follows from the very definition of TUIs given in the influential *Emerging Frameworks* paper (Ullmer and Ishii, 2001): "the physical state of interface artifacts partially embodies the digital state of the system". The advantages of such spatial syntax are manifold, but the primary intention has remained the same since the original *Tangible Bits* paper (Ishii and Ullmer, 1997): "to take advantage of natural physical affordances (Norman, 1988) to achieve a heightened legibility and seamlessness of interaction between people and information". Yet despite these advantages of spatial syntax, which have fuelled the successes of TUIs to date, the use of spatial syntax comes with disadvantages too. As a designer, the decision is not just about which form of spatial syntax to use, but whether to use it at all; as with any design decision, it is the trade-offs that are most important.

The Problem with Spatial Syntax

Tangible Correlates analysis (see Section 3.4.1) reveals the following usability characteristics of spatial syntax:

1. The use of spatial representations to model inherently spatial problem domains can achieve higher `structural correspondence`$^{\texttt{[closeness of mapping]}}_{\texttt{<TC>}}$ than non-spatial representations.

2. The use of physical representations to model inherently physical problem domains can achieve higher purposeful $\texttt{affordances}^{[\texttt{role expressiveness}]}_{<\texttt{TC}>}$ than non-physical representations.
3. Meaning is only produced through the mutual configuration of multiple objects – all meaningful representations are therefore composed of multiple objects, resulting in higher $\texttt{bulkiness}^{[\texttt{diffuseness}]}_{<\texttt{TC}>}$ than atomic (single object) representations.
4. The use of composite structures to represent information binds each component part into a structure – these structures can only persist for as long as their component parts are not required with greater importance elsewhere, resulting in lower $\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$ than atomic representations.
5. The spatial relationship between objects is generally easier to accidentally change than the internal physical configuration of individual objects, resulting in higher $\texttt{shakiness}^{[\texttt{error proneness}]}_{<\texttt{TC}>}$ than atomic representations.

This analysis suggests that whilst spatial syntax is highly appropriate for problem domains that are inherently spatial – such as the building layout in Urp (Underkoffler and Ishii, 1999) – or inherently physical – such as the landscaping TUI of Illuminating Clay (Piper et al., 2002) – there are certain aspects of spatial syntax usability that could be detrimental in certain activity contexts. In particular, the use of relational mappings to represent abstract relationships is strongly affected, since it is subject to the disadvantages of TUIs without being able to exploit familiarity with known uses and interactions of physical objects. In these situations, spatial syntax captures the use of space as part of the formal notation, reducing opportunities for spatial $\texttt{secondary notation}_{<\texttt{CD}>}$.

A compounding trade-off that exists within the different forms of relational mapping is between $\texttt{rigidity}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$ (resistance to change in configuration), and $\texttt{rootedness}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$ (resistance to change in location). In his dissertation on *The Language of Graphics*, Engelhardt (2002) presents the six graphical forms of spatial syntactic relation, each having their foundation in the Gestalt principles of visual perception: spatial clustering (from proximity); separation by a separator (using lack of symmetry to draw a distinction); line-up (from good continuation); linking by a connector (also from good continuation); containment by a container (using the notion of containing area); and superimposition (from closure). In the tangible domain, all of these relationships can be dynamically manipulated, with containment, linking and superimposition (stacking) taking on added significance due to the richness of their physical expression. Such object-to-object relations can express information regarding association, disassociation and order (ibid.). My analysis of the usability properties of these syntactic relations, using the Tangible Correlates, highlighted two distinct subtypes of complementary usability. The relations of stacking, connection, and containment are all based on *physical bonding* through gravity, linkage, and common enclosure respectively, making them easier to move and relocate as a unit, but more difficult to reconfigure due to the requisite breaking and making of such bonds. In contrast, the relations of line-up, clustering, and separation are all based on *perceptual arrangement*, making them easier to reconfigure but more difficult to move and relocate as a unit.

With respect to this analysis, the hitherto automatic use of spatial syntax in TUIs, and the definition of TUIs as embodying such spatial syntax, should be questioned. Whilst pure-spatial and constructive TUIs have enjoyed success in specialised domains with a clear and intrinsic spatial syntax, the success of relational TUIs has been more limited. This is because physical familiarity is not so clearly exploited, and the disadvantages are all the more apparent when compared with the natural alternative for manipulating abstract relational structures – a traditional Graphical User Interface. This is not to say that relational approaches to TUIs should be abandoned; rather, more time and effort should be dedicated to discovering how tangibility can be best applied to abstract domains. Given the analysis presented here, this could reasonably entail a relaxation of the need for spatial syntax, and in turn open up new areas of TUI research and application.

The existence and nature of spatial syntax in TUI design can therefore be seen as a decision that entails a variety of trade-offs, all of which must be considered in relation to the activities to be supported. Fortunately, there are a number of examples from the literature which detail the many and varied ways in which people adapt their environment to suit their personal style of work, specifically through the use and appropriation of regular physical artefacts, whose spatial configuration provides benefits to their user without the need for computational interpretation.

4.1.3 Support from Ethnographic Studies

Two prominent examples of users adapting their physical environment to suit their needs are the use of paper flight strips in Air Traffic Control, and the use of paper in office environments. The following sections will draw parallels between these uses of paper and the potential use of digitally-augmented physical tokens in an office environment.

Flight strips in Air Traffic Control

The use of paper flight progress strips in Air Traffic Control (ATC) is an often cited example of the utility of external physical representations, whose subtle affordances for coordinated work in a team environment have resulted in a history of controller resistance to the imposition of wholly-computerized systems. Ethnographic studies of ATC have been conducted by Harper et al. (1991), Bentley et al. (1992), and MacKay (1999), each of which details the context provided by an ATC room, as well as the use of paper flight strips as the underlying mechanism both of control and of coordination between "individuals [that] are individuals-in-a-team" (Bentley et al., 1992). The fact that much office-based work is conducted within similar teams, albeit in a less safety-critical context, highlighted the possibility of borrowing the concept of tangible coordinating mechanisms from ATC and applying them in the office environment. Table 4.1 pairs quotes from *Is Paper Safer? The Role of Paper Flight Strips in Air Traffic Control* by Wendy Mackay (1999), with short explanations of how these affordances could be applied analogously in the context of coordinated desk work in the office environment.

Observation from Air Traffic Control	**Design Opportunity for Office Context**
The physical layout of the strips [...] provides a temporal as well as a spatial framework for managing activity. [...] Controllers "insert" actions to do in the future.	Physical tokens could be used for the planning and tracking of generic work tasks, with temporal interpretation of linear spatial relationships.
The current strip setup reduces the controller's mental load, [...] the information is always instantly accessible in front of them.	Physical tokens distributed throughout the desktop environment would make information visible as well as physically accessible.
[The strips are part of] the controller's mental representation, helping him or her handle more information and successfully deal with interruptions.	The current configuration of physical tokens could embody the current work context, using the desktop as an external, persistent memory.
[Mental representations are offloaded] through annotations, juxtapositioning of related strips, and sliding strips to the side in their holders.	The flexibility-in-use of physical tokens would support multiple working styles through the organic development of physical organization strategies.
[The act of inserting a new flight strip into the appropriate holder] forces the controller to mentally register the new flight.	The receipt of a new token would force users to mentally register the new task, and act as a persistent reminder to incorporate it into their work plan.
Controllers must physically pick up each strip and place it somewhere; the location determines who will handle it next.	Physical task tokens could embody the responsibility to work on those tasks; exchange of tokens would entail a concrete transfer of responsibility.
Controllers often take strips in their hands as a concrete reminder to deal with that strip next.	The manual handling of tokens could help users to concentrate their thoughts on the tasks represented by the tokens.
Controllers periodically reorder the strips [which gives them] the sense of "owning" the aircraft and reinforces their memory of the current situation.	Periodically glancing at tokens and rearranging them could help users to maintain and refresh their mental model of the current work situation.
Student controllers can be observed "thinking out loud with their hands" as they touch each individual strip involved in a conflict.	Physically touching and manipulating task tokens could help users to spatially index their thought processes and help with task comparisons.
Controllers [...] actively or passively push or pull information back and forth between their periphery and focus of attention.	Users could seamlessly move physical tokens in and out of their attentional focus as required, according to the dynamic needs of their activity.

Table 4.1: Design Opportunities based on Air Traffic Control Observations

Taken together, these mappings from ATC to the office domain provide a *task-oriented* view of the role that tangible tokens could play in a generic office environment. The fundamental lesson to be learnt from paper flight strips is that users, both individually and collectively, should be given the freedom to develop their own strategies for managing the coordination of work plans, mental models, attention, and responsibilities, through the annotation, arrangement, handling, and transfer of physical tokens.

In order to participate in the many activities described here, such tokens cannot be anchored to any particular interactive surface, constructive assembly or token+constraint system – they must be free to move around from hand to hand, user to user, and desk to desk in the same fluid manner that paper flight strips are managed in an ATC centre. Tokens must therefore be meaningful in their own right, in the absence of any coupled digital representations; when they are manipulated in the context of a TUI, its role should be analogous to the radar system of ATC, which complements the provisional, plan-oriented paper flight strips by providing dynamic visual feedback about the real-time work situation resulting from their decisions (i.e. the positions of aircraft in 3D space). Such an arrangement exploits the relative advantages of physical and digital media, in particular the ability of human users to make decisions based on reasoning with external artefacts, and the ability of computers to accurately track the consequences of those decisions and feed them back to the users for consideration in future decision making processes. It is this complementary, loosely-coupled relationship between tokens and TUIs that defines the task-oriented application of tangibles inspired by ATC.

Paper Use in Office Environments

Whereas paper flight strips provide a notation for distributed decision making in ATC, in an office environment, paper documents provide a less formalised, less time-pressured means of achieving similar goals in terms of tracking decisions, coordinating resources, focusing attention, and so on. However, paper documents are used on much coarser spatiotemporal scales than paper flight strips, maintaining their relevance for longer, often persisting past their lifetime of relevance, and requiring complex spatial organisation strategies to manage the large volumes of paper-based information that result.

Seminal studies on the role of paper in organizations are presented in Sellen and Harper's (2003) book entitled *The Myth of the Paperless Office*. In the book, they give four uses of paper:

1. A*s a tool for managing and coordinating action among co-workers in a shared environment.*
2. *As a medium for information gathering and exchange.*
3. *As an artefact in support of discussion.*
4. *As a means of archiving information for groups of co-workers.*

Each of these uses is built on the duality of paper documents as both a physical indication of the existence of some information, and the physical manifestation of that information in terms of the marks on the paper. However, these roles can feasibly be decoupled. Provided that the physical indication is somehow identifiable and can in some way be used to access the underlying information, it can act as a proxy for that information in all four of the uses of paper listed above. If this decoupling between physical representation and information is digitally mediated – i.e. the physical token can be using to access digital information – then many of the advantages of physical media will be maintained, whilst avoiding a number of the problems associated with physicality. For instance, Sellen and Harper (ibid.) list the following interactional costs associated with paper documents:

1. *Paper must be used locally and cannot (without supporting technology) be remotely accessed.*
2. *Paper occupies physical space and thus requires space for its use and storage.*
3. *Paper requires physical delivery.*
4. *A single paper document can be used by only one person at a time.*
5. *Paper documents cannot be easily revised, reformatted, and incorporated into other documents.*
6. *Paper documents cannot be easily replicated.*
7. *Paper documents, on their own, can only be used for the display of static, visual markings.*

By adopting a *document-oriented* view of the role that tangibles could play in the generic office environment, in contrast to the task-oriented view inspired by ATC, many of these costs could be mitigated. Suitable token-based access to digital information would eliminate all of them except (2) and (3), which when considered in light of the ATC studies discussed previously, are actually essential parts of the subtle processes of coordination and awareness that constitute the foundation of a cohesive team.

Accepting that digitally-augmented, physical tokens share many of the characteristics of paper documents allows a number of other benefits to be realised. Sellen and Harper (ibid.) observed that paper documents were often arranged on desktops as ways of “temporarily marshalling, organising, displaying and giving access to information, and as a way of reminding workers of jobs that needed to be done or things that needed attending to”. In particular, they found the metaphor of ‘hot, warm, and cold’ documents to be helpful in understanding users’ spatial organization strategies. Hot documents were those ’on the boil’, representing information directly relevant to the current work situation. These sat on the desk to hand, ready to be carried from place to place as needed. Warm documents were those ’simmering’ – they had either just been hot, or were about to become hot in the near future. These were located close at hand either on the periphery of the desktop or in a desk drawer. Finally, cold documents were those with limited relevance to the recent or anticipated future work situation, retained ’just in case’ in a central cabinet. This strategy of using desktop space to organise information in a manner that supports workflow, marshalling resources into configurations that support the task in hand, is one for which physical tokens operating within a TUI would be particularly suited. The sensing region of the TUI could be used for ’hot’ resources, whilst the broader desktop environment could be used to structure ’warm’ resources that need to be kept within reach. Under this scheme, ’cold’ resources would be those residing purely in electronic form, not requiring the enhanced accessibility characteristics of physical tokens.

4.1.4 Identification of Problem Area

The foregoing descriptions of paper use in two distinct working environments presents a case for the important role that physical artefacts already play in working life, and motivates an investigation into how TUIs, based on physical tokens that are annotatable and spatially-reconfigurable, can support office-based individuals whose work is mutually interrelated, yet performed independently at their own desks. Combined with the two main problems arising from my initial investigative interviews in an office environment – the diminution of conversational opportunities due to the computerisation of communication, and the lack of support for auxiliary work activities – this supports a research agenda in which a TUI is both designed for this context and evaluated in it.

Such a TUI would need to address the problems of multitasking by lowering barriers to interaction, by allowing users to manage task information – completion dates, time remaining, actions to do, etc. – in parallel with the performance of those tasks. This could be achieved by giving tasks a tangible representation in terms of *task tokens*, which would be the focus of these task management interactions. In order to convey task progress to other team members, it would also be necessary to represent them in the TUI – *contact tokens* could allow users to selectively visualise one another's work load, work status, and work progress. These would help to address the problem of reduced mutual awareness between weekly progress meetings. Finally, *document tokens* could be used to take advantage of the complementary benefits of physical and digital media: the physicality of tokens would enable direct and immediate access to documents, whilst the digital nature of the underlying document would allow it to be freely edited on a remote server. Given these token types, a suitable TUI design could create new opportunities for face-to-face conversation by encouraging *task delegation* and *document sharing* through the physical passing of tokens. In this way, the tokens could be used to retrofit a layer of sociality on an increasingly faceless world of digital communication.

The general characteristics of this potential TUI solution resemble the antithesis of what I regard as the mainstream approach to TUI design:

- Rather than co-located users performing focal, collaborative activities on a single, shared tabletop interface embodying a spatial syntax, I argue for the value of a highly contrasting approach: individual users performing auxiliary, coordinated activities on their own desktop interfaces, which provide freedom of spatial use and interpretation.
- Rather than the physical exchange of artefacts resulting in diseconomies relative to instantaneous forms of digital communication, I argue that their exchange provides opportunities to bridge the gaps in understanding and awareness resulting from the lack of face-to-face conversation.

Central to the performance of auxiliary activities in parallel with primary work tasks is an approach to interaction that does not require attention, yet can focus it where necessary, and which aims to support individual users, yet conveys benefits to the whole team. This kind of interaction is exemplified by Air Traffic Controllers, who "actively or passively push or pull information back and forth between their periphery and focus of attention" (MacKay, 1999). The term I coined for this paradigm is *peripheral interaction*, which aims to capture this distinctive combination of temporal, spatial, social, and attentional characteristics:

> **Peripheral Interaction** gives users the freedom to arrange independently-meaningful, digitally-augmented physical tokens on the periphery of their workspace and away from their normal centre of attention, ready to selectively and fluidly engage those tokens in loosely related, dispersed episodes of digital, cognitive, and social use.

This definition gives rise to a number of general research questions that this dissertation will attempt to answer:

1. APPLICATION CONTEXT. What contexts, in terms of activity and environmental structures, are best suited to peripheral interaction?
2. INTERACTION STYLE. What procedures, tools and methods should we use to design for a peripheral interaction style in these contexts?
3. INTERFACE STRUCTURE. What structural forms, styles of mapping, and modes of representation can support this peripheral interaction style?
4. INFORMATION CONTENT. What are the different kinds of information that this peripheral interface structure can represent?
5. APPLICATION JUSTIFICATION. What are the essential qualities of peripheral interaction that justify its use in interacting with such information?

The first of these – the APPLICATION CONTEXT research question – has been partially addressed in the preceding part of this chapter; the suitability of peripheral interaction in other contexts will be discussed in the Conclusion (Chapter 8). The remainder of this chapter will deal with the second research question, on INTERACTION STYLE, drawing on both previous literature (Chapter 2) and my own framework for analytic design (Chapter 3). The outcome of this process will be a TUI design that supports peripheral interaction, and hence contributes to the third research question on INTERFACE STRUCTURE; this is then discussed in detail in Chapter 5 on the implementation of a peripheral TUI prototype. The final two research questions – those on INFORMATION CONTENT and APPLICATION JUSTIFICATION – will be addressed in the two evaluation chapters (Chapters 6 & 7).

4.2 Activity Analysis

The preceding sections have served to determine characteristic features of a TUI design that will fulfil the requirements of peripheral interaction in the office context. The interaction context can be described by the activity profile of suitable notations, presented in terms of Tangible Correlates (introduced in Section 3.4.1).

1. Low $\texttt{bulkiness}^{[\texttt{diffuseness}]}_{<\texttt{TC}>}$: The TUI should support a single user at their desk, complementing their existing workstation of monitor, mouse, and keyboard without impinging on their spatial resources.
2. High $\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$: Tangible representations should persist for as long as they are relevant to the prevailing work context.
3. Low $\texttt{shakiness}^{[\texttt{error proneness}]}_{<\texttt{TC}>}$: The TUI should support updates to information in a manner that is robust against accidental change, especially if changes are communicated in real-time to other team members.
4. Low $\texttt{rootedness}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$: Tangible representations should be portable, such that they may be meaningfully arranged in the desktop environment and exchanged between individuals.
5. Low $\texttt{rigidity}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$: Rapid switching should be possible between workstation and TUI, with rapid tangible interactions.

This profile is incompatible with the characteristics of spatial syntax, presented in Section 4.1.2. It suggests that, in terms of supporting auxiliary work activities at users' desks, conventional TUIs are inappropriate. However, four of these five requirements would be satisfied by a tangible interface in which individual physical tokens – sized so as to fit in the palm of the hand – are used to represent items of common interest to a team:

1. Low $\texttt{bulkiness}^{[\texttt{diffuseness}]}_{<\texttt{TC}>}$: A single token does not have a significant desktop footprint.
2. High $\texttt{permanence}^{[\texttt{visibility}]}_{<\texttt{TC}>}$: Individually meaningful tokens can persist for as long as necessary.
3. Low $\texttt{shakiness}^{[\texttt{error proneness}]}_{<\texttt{TC}>}$: Accidentally moving tokens in a non-spatial syntax has no effect.
4. Low $\texttt{rootedness}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$: Small, disc-shaped are portable and can easily be moved and exchanged.

The challenge was to create a notation based on non-spatial syntax, in which all of these Tangible Correlate criteria still held. The same notation also needed to provide low interaction $\texttt{rigidity}^{[\texttt{viscosity}]}_{\texttt{<TC>}}$ – interactions with digital information through these tokens should encounter very little resistance.

With these requirements in mind, I revisited the technology company in which I had conducted investigative interviews, with the aim of gaining some inspiration in terms of how a TUI could be integrated into the existing workspaces and work practices of potential users. My interest in fine-grained phenomena – how users transition between different activities, what they do with their hands, and how they interact with their desktop 'habitat' – suggested that contextual observation based on video recording would be an appropriate approach. An engineer volunteered to take part in this video study, and was subsequently recorded during a morning's work at his desk. From the footage, a 30-minute period was selected in which the engineer was performing a typical workstation-intensive task (software debugging). This was then subjected to "micro-analysis", in which all of the engineer's actions were coded on a per-second basis for the whole of this 30-minute period.

Figure 4.1 gives a visualisation of this video analysis, performed in an Excel spreadsheet. For each second of video time, the actions of the left and right hands were both coded according to their state: *active* (using the keyboard and mouse); *ready* (in a 'home' position on the keyboard or mouse); *idling* (any non-goal-directed behaviour such as resting, tapping, grasping and so on); and *other* (writing on the whiteboard, using the telephone, eating, drinking, adjusting chair, etc.). Changes in application focus on the workstation were also coded, as well as the subject's current activity (typing, scrolling, reading, thinking, talking, etc.), various forms of conversation participation (office conversation, telephone conversation, etc.), and sources of background noise (doors opening and closing, private conversations, etc.).

4.2.1 Opportunities for Interaction

The first analysis conducted was to examine the physical actions of the hands during the study, since patterns of utilisation and availability could demonstrate opportunities for desk-based tangible interaction. Figure 4.2 shows the results of this analysis.

For about half the time, both hands were engaged in typing or coordinated use of keyboard and mouse. For a quarter of the time, however, the subject's hands were not engaged in workstation-based activities; rather, the subject was attending to events happening around them. These including answering a ringing phone, checking to see who had just entered or left the room, and participating in office conversations. This sizeable proportion of non-workstation activity time during what was a workstation-intensive task (software debugging) presents a design opportunity, in terms of both exploiting natural breaks in activity to perform peripheral updates, and providing users with a way of recovering their work context after interruptions.

Peripheral interaction, however, is not just about having time to spare in which tangible updates can be made, but about being able to make such updates episodically according to the

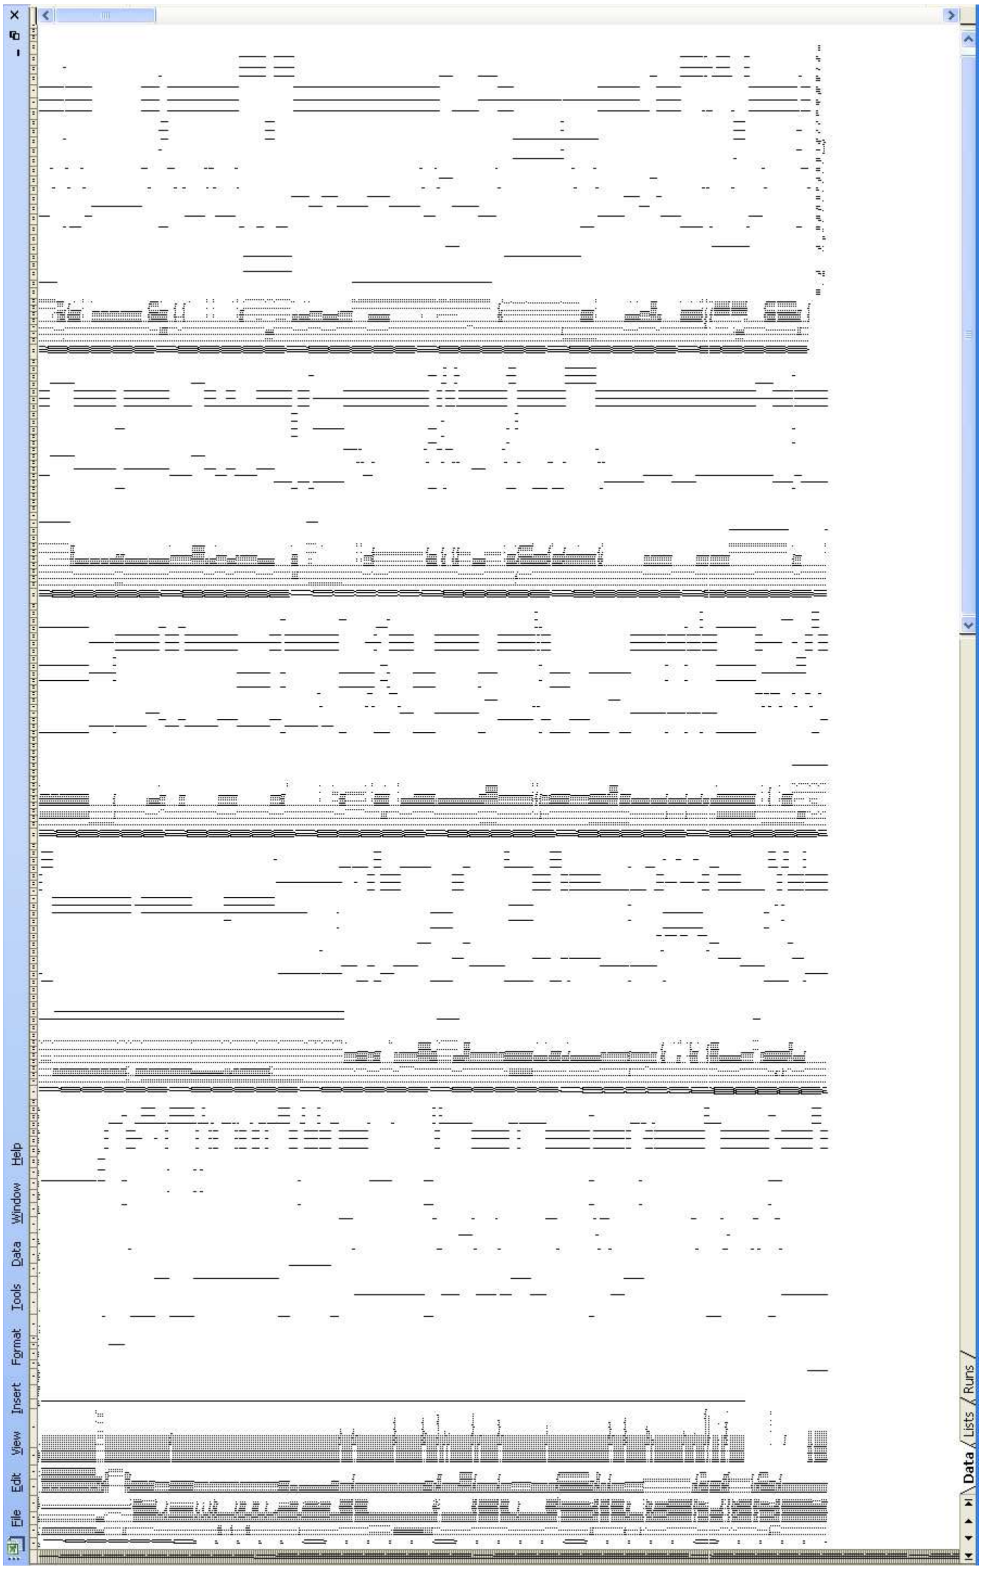

Figure 4.1: Spreadsheet-based Micro-analysis of Engineer Actions: Indicates the volume and method of data analysis (not intended to be legible).

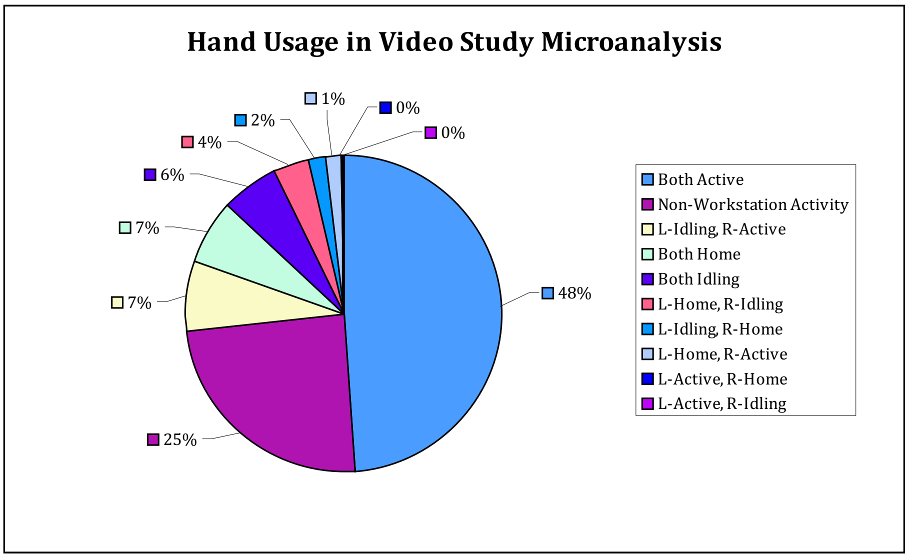

Figure 4.2: Hand Usage in Video Micro-analysis: Breakdown of independent and coordinated hand use in an 1800s video segment of an engineer performing a software debugging task.

flow of primary work activities. To examine how changes in the subject's attentional focus might influence the duration and frequency of potential peripheral interactions, the spreadsheet of user actions was used to generate a chart that highlighted the cyclical repetition of action followed by thought followed by action. For this purpose, "thought" was defined as a user state in which the focus of their gaze remains steady, and where their hands are not purposefully interacting with anything (this definition does, however, permit idling behaviours of the kind mentioned previously). The user state of "action" is the complement of this focused state of attention. Figure 4.3 presents this chart.

This visualisation depicts a distribution of thought-action cycles in which periods of thought (mean 5 seconds) are linked by periods of sustained action (mean 26 seconds), and where the longest period of sustained action is just over two minutes (140 seconds). Even in this workstation-intensive task, periods of user action are fragmented by many small periods of thought and reflection. It is in the transitions between these micro-cycles of work, lasting approximately 30 seconds each, that opportunities for peripheral interaction arise.

4.2.2 Observed Asymmetry of Bimanual Action

The analysis of hand usage (Figure 4.2) also revealed notable differences between the subject's left and right hand sides. When only one hand was actively operating the keyboard or mouse, the right (dominant) hand was employed for twice the duration of the left hand (164s vs. 81s); this asymmetry was reversed when idling (non-workstation) behaviour was considered, during which the left hand was observed tapping, fiddling, scratching, face touching,

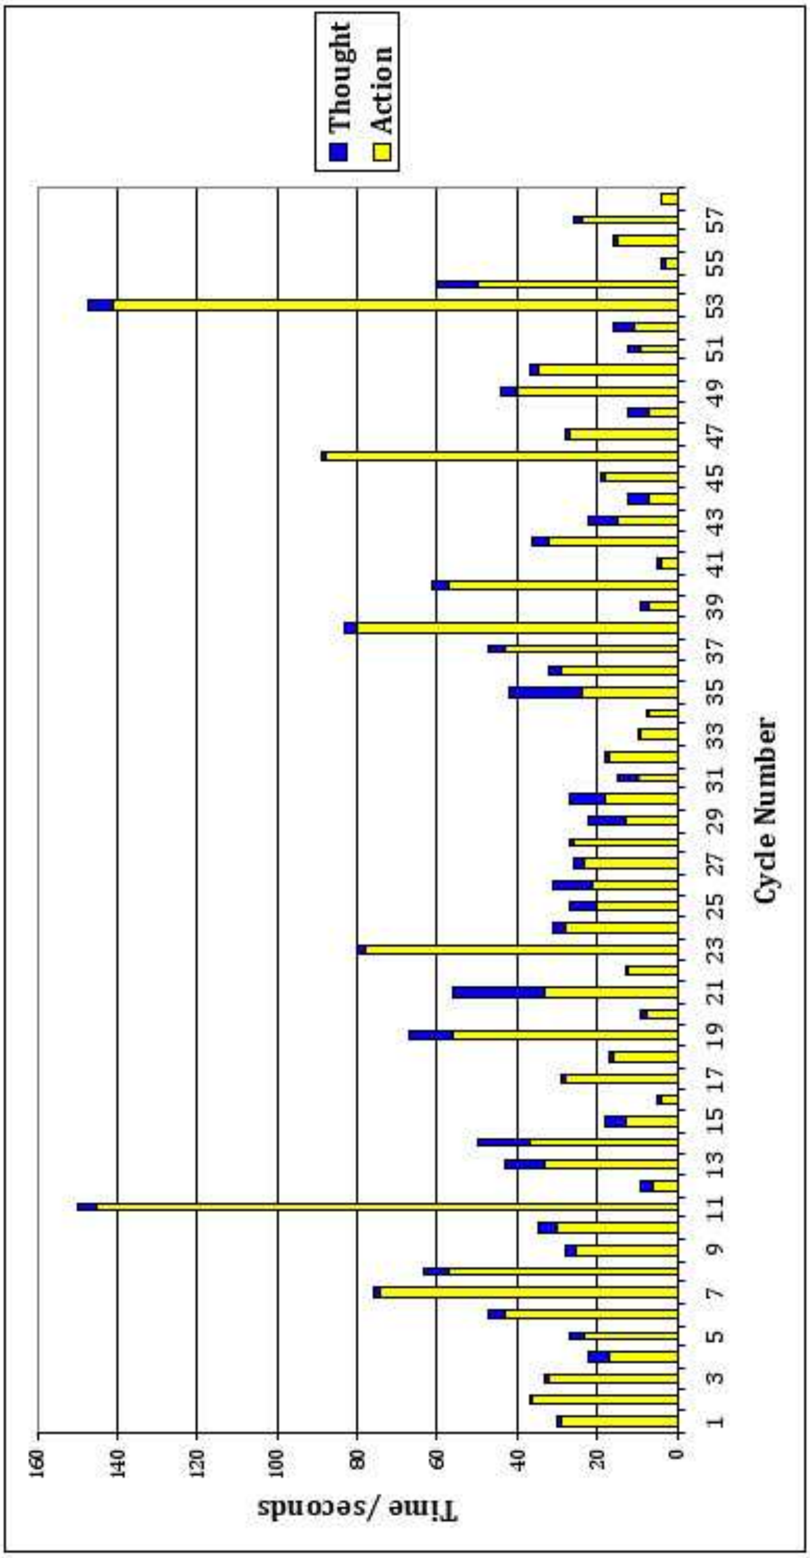

Figure 4.3: Cycles of Thought and Action in Video Study: Each bar indicates a single cycle of a period of action (yellow) followed by a period of thought (blue). The heights of the yellow and blue bars represent the respective durations of these periods, and subsequent bars represent subsequent cycles, which repeat until the end of the 30 minute analysis period.

screen grasping, pointing, gesturing, and resting for almost four times the duration of the right hand (283s vs. 71s).

When not using the keyboard and mouse, the engineer regularly used his right hand to touch objects to the right of his keyboard, without diverting his visual attention from his monitor screen – his desk phone, mobile phone, mug of coffee, and bag of confectionery all being located, grasped, and utilized in this eyes-free manner. In contrast, he did not touch the various piles of documents and books on the left hand side of his workspace, although he did refer to and update the whiteboard located above them throughout the recorded period.

These measurements represent a marked asymmetry in terms of hand usage, as well as lateral specialization of the workspace either side of the workstation. They support the concept of a TUI structure that takes advantage of this lateral specialization, whilst providing interaction opportunities for the under-utilized left (non-dominant) hand. Any TUI exploiting this design opportunity needs to take into account Guiard's (1987) work on the *Asymmetric Division of Labor in Human Skilled Bimanual Action*, which abstracts away from the biomechanical and physiological complexity of skilled, coordinated bimanual activities, modelling the hands as abstract motors that tend to operate in series. This is the basis of Guiard's *kinematic chain* model, which accounts for the following phenomena commonly observed in the diverse activities of handwriting, violin playing, golf swing performance, sewing, driving a screw, and others:

1. *Right-to-left spatial reference in manual motion*
2. *Left-right contrast in the spatial-temporal scale of motion*
3. *Left-hand precedence in action*

Guiard (ibid.) summarises these three characteristic features of skilled bimanual action: "in general, right-hand motion is built relative to left-hand motion, corresponds to a temporal-spatial scale that is comparatively micrometric, and intervenes later in the course of bimanual action". These characteristics account for our everyday actions in the real world, in which bimanual cooperation entails both hands operating on the same physical objects, in the same physical space. However, in terms of interacting with computers and other electronic devices, bimanual cooperation is not constrained by the limits of the natural physical world. Interaction with virtual objects tends to be through parallel, rather than serial, composition of the hands. In typing, this parallel cooperation is symmetric, whereas in the coordinated use of the keyboard and mouse – such as the navigation of virtual three-dimensional worlds – this parallel cooperation is asymmetric. The standard control mapping in such circumstances is for the left hand to operate the keyboard for categorical control of movement, option selection, etc., and for the right hand to operate the mouse for continuous control of view direction. Interestingly, the operation of these controls exhibits the same three characteristics listed above, but with parallel rather than serial cooperation between the hands:

1. Changes of view direction with the right hand are relative to the virtual spatial frame of reference set by the left hand, which controls movement in the virtual world.

2. The discrete keypress operations performed with the left hand are on a coarser spatiotemporal scale than the continuous mouse-pointing actions of the right hand.
3. Option selection and movement via left hand key-presses precede right hand pointing actions.

This observation suggests that our interaction with virtual worlds may be influenced by our asymmetric bimanual conditioning in the physical world, which in turn is influenced by the neurophysiological differences that lead to lateral manual preference. As such, it is possible to imagine a *conceptual kinematic chain* that links the virtual output of the left hand to the virtual input of the right hand, despite their parallel, decoupled relationship in the physical world.

It is therefore reasonable to construct an asymmetric, bimanual interface based on this concept, in which the two hands operate in physically decoupled spaces, performing conceptually coupled operations on objects with coincident physical and virtual representational forms. My interpretation of this is a token-based tangible user interface, in which there is tight coupling between the physical and digital representations, but loose coupling between these representations and their control. This division of labour within the interface should correspond to the natural division of labour that occurs in skilled bimanual interaction.

Assuming tokens represent digital information objects with multiple attributes (for a right-handed user):

1. the left hand should lead, being used for coarse-grained selection between multiple tokens and their various attributes, and setting a spatial and attentional frame of reference;
2. the right hand should follow, acting within this frame of reference, performing fine-grained manipulations of the selected attribute with a single control device.

By mirroring the cooperation between the time-multiplexed operation of the mouse, and the space-multiplexed operation of the keyboard, the resulting interface structure naturally complements the traditional workstation setup of monitor, mouse and keyboard.

This abstract interface structure can be transformed into a concrete TUI design by answering the following questions:

1. How should the TUI use physical tokens to represent digital information objects with multiple attributes?
2. How should the TUI support coarse-grained attribute selection through left-hand operations on tokens?
3. How should the TUI support fine-grained attribute manipulation through right-hand operation of a control device?

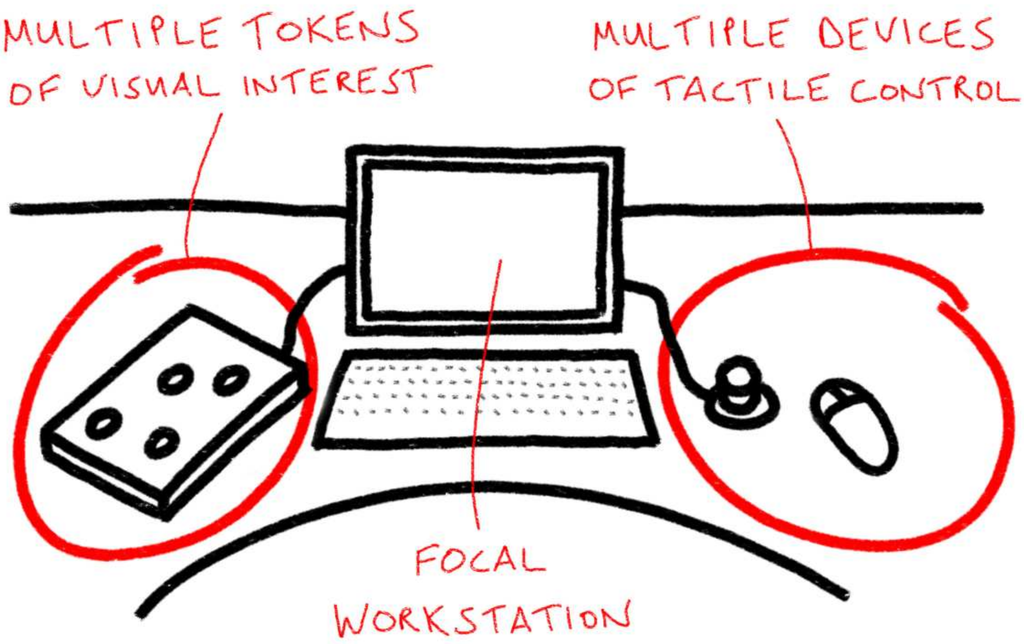

Figure 4.4: Design Sketch of a Bimanual TUI for Peripheral Interaction

4.2.3 Asymmetric Interface Design

The TUI design I use to support peripheral interaction is the result of answering the questions posed in the preceding analysis, and is shown in Figure 4.4:

The primary component of the interface is an interactive surface placed to the left of the keyboard (for a right handed user). When a physical token is placed on this surface, its position and identity are determined, and the information attributes associated with that token are displayed underneath and around it as a *halo* on the interactive display surface. Each attribute of a token is selected by *nudging* the token in the direction of the attribute as it appears within the surrounding halo. Regarding the number of attributes such halos should represent, there is a trade-off between increased information and increased precision of selection. Given that the TUI is designed to support peripheral actions, the number of attributes should be limited to what can be perceived 'at a glance', and which can be selected between using a small set of coarse-grained 'nudge' gestures. Restricting nudge gestures to the four axis-aligned directions relative to the interactive display surface provides a reasonable balance between the information content and the required levels of perceptual acuity and motor precision.

The final design decision as to the nature of the control device draws on the Tangible Correlate analysis of continuous value representation and control presented in Appendix A.1. The nature of information attributes that can be rapidly perceived and manipulated will typically be of a single dimension (quantities, indices, lists, etc.) and so a single-degree-of-freedom (1-DoF) control device is required. Of the six prototypical 1-DoF device structures – knob, joint, length-slider, position-slider, length-screw and position-screw – only the *knob* has the required characteristic of 'statelessness' that allows it to be multiplexed between many information attributes[1]. This characteristic also allows the knob to be acquired and operated in an eyes-free manner, with its position between the keyboard and the mouse minimising the average 'homing' distance required, and exploiting the complementary muscle memory and spatial memory arising from many interactions with a device in a fixed location.

In the resulting interface, the objects of representation – the physical tokens – are decoupled from their means of control – the single control knob. Actions on the tokens alone cannot change the digital state of the interface – it is the coordinated bimanual displacement and control device operation for individual tokens that modifies the underlying digital information. The arrangement of tokens is therefore not in itself meaningful – moving tokens does not change the digital state of the interface. The interface is *free of spatial syntax* and the problems associated with it, which were presented in Section 4.1.2.

The deliberate recruitment of both hands also ensures that actions are intentional. This helps to guard against accidental change, since the extent of change possible by accidentally knocking tokens is a change in the identity of the currently selected attribute, rather than a change in its value. This 'bimanual safeguard' allows users to make rapid, intentional changes whilst sustaining tokens being moved around on the surface, taken off and even passed between individuals for use on other surfaces. The combination of complementing the existing workstation and supporting fast, robust interactions, means that rapid switching is possible between the workstation used for primary work tasks, and the TUI for auxiliary work activities. In summary, the use of a non-spatial syntax based on coordinated bimanual interaction with representational tokens and a single decoupled control device, meets all of the interaction requirements of the usage context.

4.3 Mapping Analysis

The TUI design resulting from further context analysis and supported by activity analysis will now be described in terms of its physical–digital mappings. The decoupling of representation and control results in a style of mapping quite unique within published work on TUIs, yet offers a wealth of interactional advantages in its intended context of use.

The decision not to incorporate a spatial syntax into the TUI design means that the TUI does not fall into a conventional style of mapping – pure-spatial, relational, or constructive

[1]Rotary devices are the only kind that can be infinitely manipulated along their 1-DoF – slider and screw devices both have bounded ranges that are only appropriate when tightly coupled to a single information attribute. Of the two prototypical rotary devices, the angle of a joint-based device exhibits a representation of state that is misleading once the control is mapped to a different attribute. On the other hand, a free-turning knob without markings has neither limits to its range nor any perceivable state, making it the only suitable choice as a pure 1-DoF control suitable for multiplexing between many information attributes.

(Ullmer and Ishii, 2000) – nor does it conform to the MCRpd structure of TUIs (ibid). The design can best be described as having a *decoupled control structure* employing a *temporal syntax*, resulting in an *asymmetric bimanual interface* (Figure 4.4).

In terms of action mappings, this interface archetype exhibits low *representational indirection* between the locations of physical tokens and their digital halos, but high *control indirection* between the locations of the single knob and the many tokens. This control indirection is embraced as a necessary part of the conceptual coupling between the two hands, spaces and objects discussed in Section 4.2.2. Compatibility is high is both cases: there is high *representational compatibility* between the movement of tokens and their halos, and high *control compatibility* between physical knob actions and visual attribute changes.

The attribute mappings of the interface archetype are similarly of two kinds. It has low *representational integration* since the many possible nudge directions map onto a small quantised number of attribute selections (typically four), but high *control integration* since the 1-DoF of the control knob maps onto one-dimensional attributes. This low representational integration is a necessary part of the bimanual control schema, enabling coarse-grained attribute selection. Regarding coupling, there is high *representational coupling* between tokens and their digital augmentation over time (high space-multiplexing of representation), but low *control coupling* between the single knob and the multiple attributes of the many tokens (high time-multiplexing of control). This 'decoupling' of representation and control is the very essence of this unique style of mapping.

Finally, the TUI design does not have any *temporal mappings* that determine future behaviour in a programmatic way, although there is nothing in the structure of the interface archetype that would preclude such functionality.

4.4 Meaning Analysis

Whereas mapping analysis examines the deep structure of the interface design, meaning analysis is about how users will perceive and interpret its surface 'look and feel'.

Perhaps the most important aspect of this analysis is the relative sizes of the tokens and interactive surface, and how the user will interpret the intended use of the TUI based on this. I decided to use poker chip-sized tokens – discs of approximately 35mm diameter and 4mm depth – to provide a balance between users being able to visually identify tokens from their surface markings, and being able to manage many such tokens spread throughout their desktop environment. The affordances of chip-sized tokens for this are manifold: lining tokens up, stacking them, placing them in containers and so on, help to both organise knowledge using tokens, and to organise those tokens using the existing physical environment. Such tokens can be transported easily in the hand or pocket for exchange or discussion with other people.

The size of the interactive surface relative to the size of tokens provides another important affordance. Given that the two dimensions of the interactive surface are not meaningful in themselves – they are simply used to provide an area in which the digital attributes of tokens can be visualised and selected – the size of the interactive surface can be chosen according to the trade-off between loss of desk space and gain in visualisation space. A surface size that can comfortably accommodate a handful of poker chip-sized tokens is likely to require an area of approximately an A4 piece of paper, which should not present a significant burden to users. This would also keep the desktop footprint of the interactive surface relatively small and encourage use of the surface as a focus for only those 'hot' work activities that are immediately relevant – merely 'warm' items of interest should then be stored in meaningful places in the desktop environment, whilst 'cold' items should have their tokens recycled for future reuse.

The design of a TUI's visual meaning should not proceed in isolation from the design of its haptic meaning. In my TUI, the choice of Perspex as the token material provided tokens with sufficient weight that purposeful nudges are required in order to displace them on the surface (i.e. to select attributes), but of insufficient weight for the tokens to damage the display surface if casually thrown onto it. Similar consideration regarding the control knob resulted in the selection of a Griffin Powermate USB controller, which combines the physical characteristics of low friction and low rotational inertia desirable for rapid yet accurate manipulations.

Tokens can also provide an outlet for idle 'fiddling' behaviour; I therefore decided to make the tokens haptically-aesthetic – pleasing to touch, heft, hold and stroke. A simple way in which to do this without compromising the 'markable' surface area used for human and computer identification of tokens is to manufacture them such that the circumference of each token has a distinctive texture. Another way to do this is to encourage users to attach haptically-distinctive materials and objects to tokens, by embedding a circular recess in the face of each token and supplying a range of such physical adornments with the TUI. I have used both of these approaches in the implementation of the TUI, to be discussed in the next chapter.

From a semiotic perspective, the physical tokens symbolically stand for digital information. This link is arbitrary and needs to be learned by convention, but is both simple and memorable. However, the link between any particular token and its 'content' is not purely symbolic: it is denoted through surface annotations of the token that literally describe its digital information content. The makes identification easier once the token is in hand, but not at a glance. For this purpose, the attachment of physical materials or objects to tokens can aid the recall of token identity as well as its recognition in the environment, assisted further if the choice of attachment is based on a self-constructed metaphoric correspondence.

Relations within and between tokens and the desktop environment, although not interpreted computationally by the TUI, can nevertheless provide a physical means through which users can furnish their environment with meaning. Tokens can be stacked, laid out in lines or clusters, placed in indexical relation to already meaningful desktop objects, or organised using physical props acquired or constructed by the user specifically for the purpose of token management and presentation. The TUI is accompanied by a large construction set of Technic-Lego with which to build structures (racks, containers, etc.) for this very purpose. Note that these concerns for the use of tokens off the surface are relatively novel: most previous TUIs were simply demonstration prototypes that were not designed for multi-session use in real contexts with live data, and can therefore rely purely on digital representations to distinguish the identity of tangibles within a single session of use.

Appendix A.2 shows how this bimanual scheme of interaction can also be explained in terms of Lakoff's (1987) image schemata and metaphorical mappings.

4.5 Appropriation Analysis

The final stage of applying the analytic design process to my context of choice was an analysis of how the TUI design might impact upon users' work practices through a process of appropriation. The questions used to probe potential appropriations of both cognitive and social nature are taken from Section 3.7 and discussed in turn here:

How does the TUI support cognitive externalisation *– in what ways might tangibles be used as prospective memory aids, or reminders?*

Task and document tokens can be used as *indexical reminders*, arranged in meaningful places – next to the phone, on the seat of the chair, along the top of the keyboard, in the wallet, etc. – to remind users both what to do, and when to do it. The small but distinctive form factor means that many such tokens can be employed as reminders, and that they will stand out as salient features of the environment.

How does the TUI support cognitive simplification *– in what ways might tangibles be used to solve problems in the world, rather than the head?*

The pre-digital scheduling of tasks can be done 'out of band' by rearranging task tokens on the desktop or meeting table. These can be instantiated as digital representations when they move into the time horizon for workload planning, but remain as a provisional, offline plan until that point in time – manifest in a line-up of task tokens in a dedicated 'medium term planning' area of the environment, such as the edge of a nearby window sill. Tokens can act as *cognitive anchors* for their respective digital content, allowing inter-content relations to be explored freely in the world.

How does the TUI support cognitive tracing *– in what ways might tangibles be annotated or marshalled into and between meaningful configurations?*

Users can evolve their own personal schemes of token annotation, for example annotating the face of document tokens with the initials of team-members who share access to that document. The same applies to users' desktop environments, onto which users can project their own personal meaning in terms of place. For instance, the general proximity of a token to the surface might be used consciously as a measure of 'warmness' – the closer a token to the surface, the more likely it is to be used again in the near future. Alternatively, tokens could be categorised into levels of 'importance', for instance by placing them near to existing desktop objects of varying value, e.g. imagine a discrete scale increasing in importance from the pad of paper, to the desk phone, to the keyboard, to the monitor, to the surface itself. The possibilities are many and varied, and the more investment users make in their personal systems of meaning, the longer those relationships will remain meaningful in their richly expressive environment. Sanders and McCormick (1987) describe a variety of strategies for the spatial arrangement of regular physical devices in the workplace, including arrangement by *importance*, *function*, *frequency-of-use*, and *sequence-of-use*. The arrangement of tokens can piggyback on these existing strategies to leverage established systems for the external creation of meaning.

How does the TUI affect social actions towards things – *in what ways might people act towards tangibles based on their functional but also their symbolic social meaning?*

For the implementation of task tokens (to be discussed in Chapter 5), I decided to use token edge texture as a symbolic indicator of the team member who *owns* that particular task. This may or may not be the person who set up the task with a name and initial set of attributes; nevertheless, the token and its associated content 'belong' to the token owner, and should eventually be returned to them. Pressure to do so is increased by limiting the number of task tokens assigned to each team member to twenty, with the aim of rendering them a scarce resource and hoping that the 'guilt factor' will encourage team members to return others' task tokens once the tasks have been completed. This socially symbolic distinction between the ownership and possession of task tokens might entail different attitudes towards task completion.

How does the TUI affect social production of meaning – *in what ways might tangibles be used to create new kinds of meaning in ways that didn't exist or weren't apparent before?*

Tokens might be arranged in a manner such that they are not just presenting their existence to the user, but to the team and its wider community (who may not have their own TUI). They might be displayed as 'badges' that indicate certain aspects of users' work roles, such as rights and responsibilities. The exchange of tokens between team members can also create new meaning, through the relationships of agency and expectation that respectively result from the token mediated transfer of rights and responsibilities.

How does the TUI affect social interpretation of things – *in what ways might individuals change their interpretation of tangibles over changes in place, time, and situation?*

The meaning of tokens is based on a triadic relationship between person, token, and environment. If the token is passed to someone else, there is no guarantee that it will hold the exact same meaning for them. The best that can be achieved is a verbal communication of meaning as an adjunct to physical transfer; this is both necessary and desirable from the standpoint of increasing opportunities for face-to-face conversation. The token recipient will return to their desk at some point and integrate the token into their personal system of environmental organisation, but this is unlikely to be the same as for the previous person in possession of the token. The meaning of tokens will also change with the passage of time – the task token for yesterday's completed task can be reused as today's new task. Changing situations can also put token meaning in a state of flux – for instance, an incoming high priority, high urgency task will shift the frame of reference against which work had previously been planned; the interpretation of existing tokens will thus need to be adjusted accordingly.

To conclude this discussion of appropriation analysis, it is necessary to say that the exact nature of appropriation is highly dependent on the nature of the individuals involved. However, the above analysis has served to highlight potential nuances of social token usage that could form part of the TUI presentation to potential users, allowing them to reflect on what appropriations would be possible and desirable within their own work context. By not mentioning potential cognitive and social appropriations, there is a risk that the system will be judged purely on its interactional characteristics and undoubtedly compared to a touch screen. In my experience of presenting the concept of peripheral interaction to other people, often via demonstration of the prototype interface to be described in the following chapter, I have found that emphasising utilisation of the broader spatial and social context provides adequate justification for the choice of a tangible user interface, and greatly accelerates its acceptance as a real alternative to purely WIMP-based interaction.

4.6 Work Related to Peripheral Interface Design

Whilst the combination of bimanual tangible interaction and peripheral displays is a novel contribution to HCI, the concepts of bimanual interaction and peripheral displays have each spawned their own research areas. Selected significant works in these areas will now be presented.

4.6.1 Bimanual Interaction

In addition to Guiard's (1987) influential paper on bimanual interaction from the perspective of cognitive neuroscience, there have been many other studies of two-handed input in HCI. In the study by Buxton and Myers (1986), two experiments were performed: one based on a compound selection/positioning task, the other on a compound navigation/selection task. In the former experiment, novice subjects automatically adopted a bimanual scheme of input, operating the selection and positioning transducers in parallel, with task completion time strongly correlating with the degree of parallelism employed. In the latter task, the two handed input condition significantly outperformed the single handed input condition.

The experimental conditions were broadened in the study by Kabbash et al. (1994) to include two "asymmetric dependent" modes of bimanual interaction in which the action of the right hand depends on the action of the left hand, as in my bimanual interface. The four conditions of a colourised 'connect-the-dots' task were: 1) right-tearoff menu, a one-handed technique; 2) left-tearoff menu, a symmetric two handed, two cursor technique; 3) palette menu, an asymmetric technique where the left hand controls menu position and the right hand controls item selection; and 4) toolglass menu, an asymmetric technique using a transparent palette menu enabling simultaneous selection and application of colour by 'clicking-through' the menu onto the underlying object. The overwhelmingly superior technique was found to be the asymmetric toolglass menu, although the other asymmetric technique – the palette menu – was found to be the worst. The authors point out that this since this cannot be due to lack of motor skills, the difference is likely to be cognitive in origin: whereas with the toolglass there was only one correct location, i.e. directly above the target object, users were often unsure where to place the palette menu and did not fully appreciate the real-world analogy to the artist's palette.

In later experimental work by Hinckley et al. (1997), a finding of particular relevance to my proposed scheme of bimanual interaction is that in recalling the six degree-of-freedom posture of the hands, subjects performed significantly better in the reproduction of two-handed postures than in the reproduction of single hand posture. The results suggest that the two hands together form a "hand-relative-to-hand" frame of reference, and that this bimanual frame of reference combines sufficient perceptual cues that information can be encoded in a manner that does not necessarily rely on visual feedback. The corollary for my TUI design is that the engagement of both hands in nudge-turn actions could facilitate more accurate reproduction of the user's interaction posture over time, reducing the cognitive demands of consciously remembering where the control knob and interactive surface are located, and speeding up motor performance accordingly.

Perhaps the most relevant experimental evaluation of bimanual interaction is a toolglass-only variation of the 'connect-the-dots' experiment, by Balakrishnan and Hinckley (1999). Here, the experimental conditions involved manipulations of device space – whether the two physical devices operate in the same physical space with a single fixed origin, or in disjoint physical spaces each with their own fixed origin, or in disjoint spaces with device-relative origins (like a mouse). The feedback mechanism was also manipulated – with visual feedback, the toolglass is always visible, whereas with non-visual, kinesthetic-cue-based feedback, the left-hand operated toolglass is only visible when in close proximity to the right-hand operated cursor position. The results suggest that both two-handed performance and the Guiard principle of right-to-left spatial reference are robust against variations in the kinesthetic reference frame of the two hands – i.e. the hands can operate on distinct physical objects in disjoint physical spaces and still co-operate in the performance of a common task, provided there is adequate visual feedback. These experimental findings directly support the informal argument of Section 4.2.2 for 'conceptual coupling' between the hands.

4.6.2 Peripheral Displays

A particularly influential early project relevant to work on displays for awareness was the Portholes project of Dourish and Bly (1992). This is a system designed to support awareness within distributed groups of users, by providing a shared "media space" that combines and presents concurrent audiovisual feeds from users' desks. However, technical challenges of the era meant that the Portholes viewer was run as a regular application on users' desktops, limiting its ability to support peripheral awareness. Subsequent technical advances have made multiple monitors a reality, and an early field study of multiple monitor use at Microsoft by Grudin (1999) revealed that people do not treat a second monitor as "additional space", but as a distinctive channel for the communication of a different kind of information. A common use of a second monitor is for *secondary activities* – in particular, the monitoring of the asynchronous communication channels provided by email, instant messaging, and so on. The area of Computer Supported Cooperative Work (CSCW) also has a long standing tradition of maintaining awareness within geographically distributed social groups in work settings, although this is often during interaction with focal groupware applications.

Peripheral displays are not limited to secondary monitors, and MacIntyre et al. (2001) have investigated the use of large, shared, wall-based peripheral displays that "complement existing focal work areas" and support "the natural flow of work across these two settings". This kind of peripherality is also related to the "ambient media" of the Tangible Bits vision (Ishii and Ullmer, 1997). An attempt to 'disentangle' the terminology used to describe displays that convey non-task-critical information is presented by Pousman and Stasko (2006), making a distinction between ambient displays, peripheral displays and notification systems: ambient displays are a proper subset of peripheral displays, having "pointed aesthetic goals" and presenting only a "very small number of information elements", whilst notification systems are motivated by "divided attention situations" in which information presented could potentially be consumed via a focal or a peripheral display.

Chapter 5

An Office-based Peripheral TUI

The prototype TUI implemented according to the design of the previous chapter is presented from two different perspectives. In this chapter, I will describe the general system structure and interaction with the various token types; in the next chapter, I will present some of the underlying architectural and technical details of the software implementation.

In my system, physical tokens are used to represent items of common interest within a team or work group, and interactions with the digital information associated with these tokens takes place on a personal interactive surface located to the left of the user's keyboard. Figure 5.1 provides an annotated screenshot of a typical token layout on such an interactive surface. Running along the top edge of the surface is a timeline-based calendar, which is used in conjunction with task tokens to indicate planned completion dates and to visualise the effects of manipulating the factors that impact upon workload: changing estimates of work-time remaining for a task, sharing time between 'overlapping' tasks, and changing the number of hours dedicated to task work in one or all working days. The yellow-and-red halos belong to task tokens, and provide an 'at a glance' visualisation of multiple task attributes that also permits rapid, low-attention interactions where necessary. In particular, task tokens enable users to adjust task completion dates, manage lists of actions 'to-do', re-estimate work-time remaining, and track the amount of time spent on each task so far. This dynamic form of time tracking aims to encourage timely recognition and revision of overly optimistic estimates of work-time remaining, by persistently presenting users with an indication of how their estimates compare to reality.

The simpler yellow-and-blue halo belongs to a document token, which acts as a physical 'shortcut' or 'hyperlink' to an underlying document. Document tokens also provide a means of achieving document-activity awareness through social access control – in order to access a document, a user must have its document token on their interactive surface; and to be in possession of its document token, they must have either created the document, or another user must have given a "cloned" copy of the document token to them. This constraint means that each user with a document token on their interactive surface can see which other users are similarly 'viewing' that document token, creating opportunities for ad-hoc collaborative authoring/editing by passively communicating mutual convenience. Finally, the green-and-blue halo belongs to a contact token – that is, a token that persistently represents a member

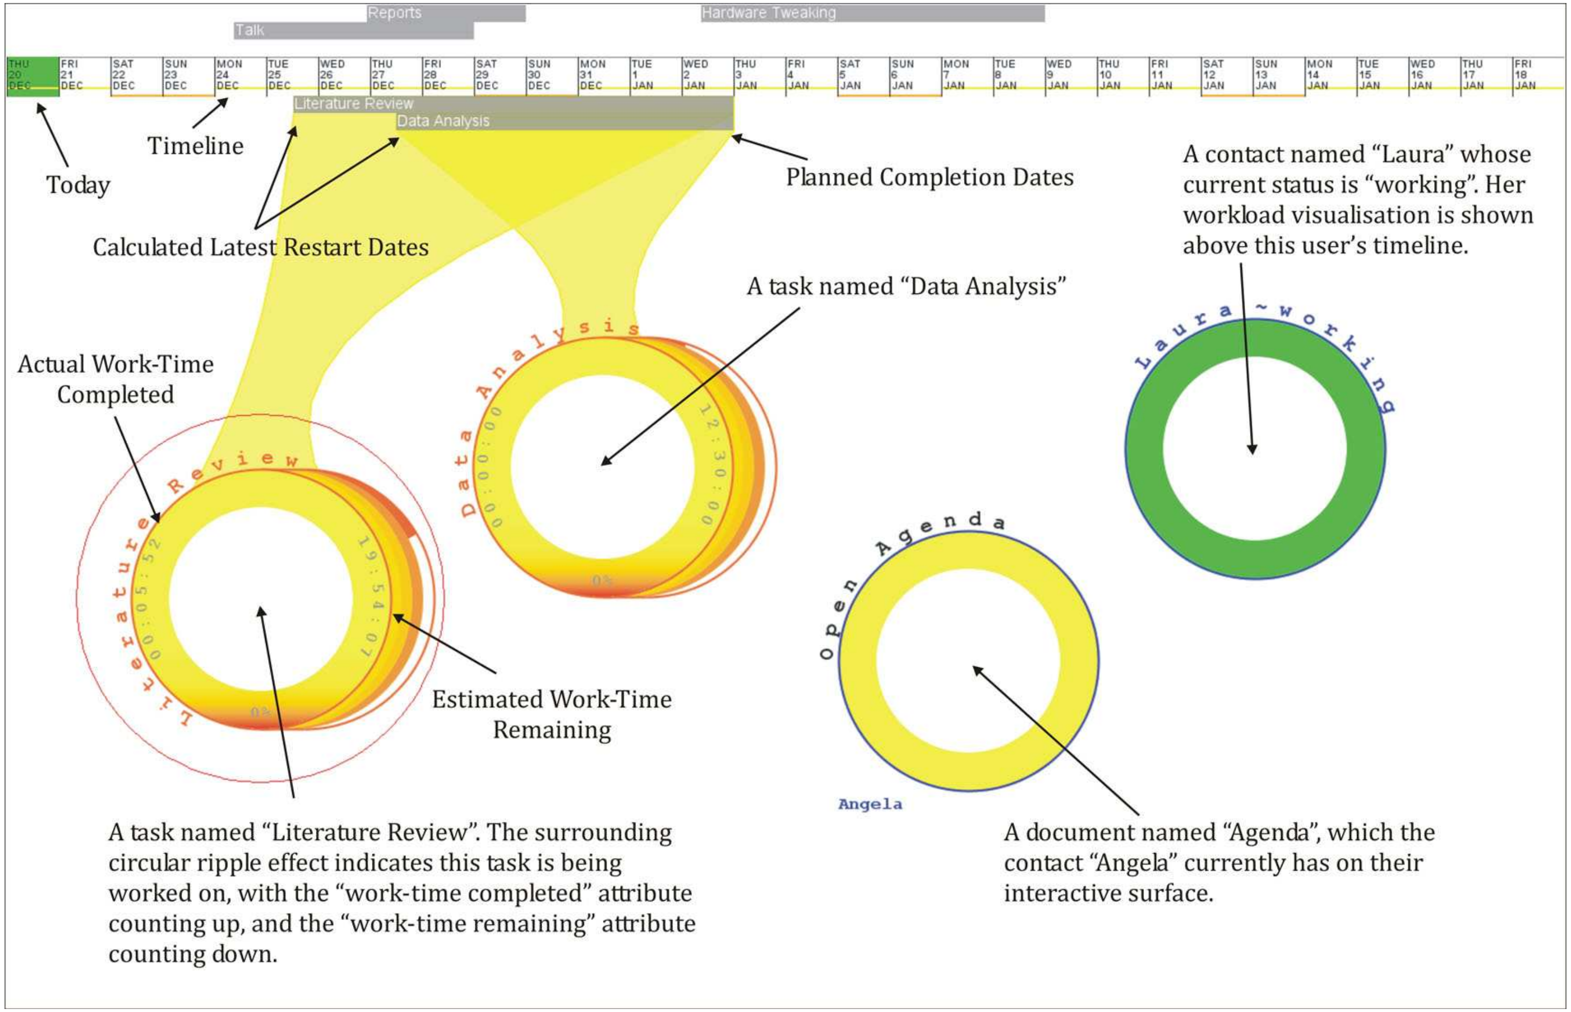

Figure 5.1: Screenshot of an Interactive Surface: A typical arrangement of token halos, captured directly from the tablet PC.

of the work group, each of whom has a TUI of their own. Each user has a "contact" token for all work group members including themselves, and any user is able to inspect the current workload of any of their "contacts" by placing the corresponding token on their interactive surface. This has the effect of inserting the contact's workload visualisation above the user's own calendar-based timeline, allowing side-by-side comparison of work schedules. This real-time 'window' onto the work of others has the potential to improve awareness in between face-to-face meetings, easing coordination between group members through the reduction of uncertainty.

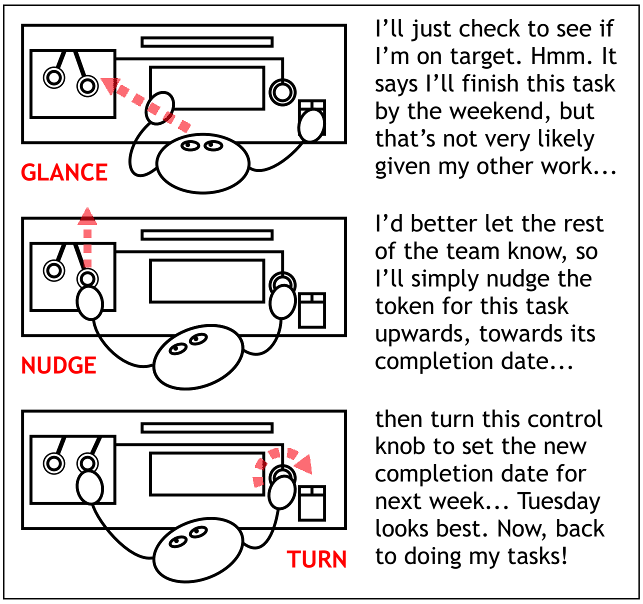

Figure 5.2: Storyboard for Peripheral Interaction

A storyboard representing the need for peripheral interaction with these token types, based on the design work from the previous chapter, is shown in Figure 5.2. The realisation of this interaction style, in which the actions of the hands are dynamically coordinated in fast, fleeting interactions with my TUI prototype, is shown in Figure 5.3 (top), as is the structure of the TUI and its integration with the existing workstation (bottom). However, the use of tokens is not solely to provide a means of direct, tangible interaction with digital information. Their physicality enables them to provide benefits beyond those conventionally delivered by traditional GUI-based technology, by reaching out from the virtual to the physical desktop. Figure 5.4 presents an illustrative appropriation of the interface in which the desktop environment itself has been rendered meaningful. The physical representation of information can provide not only cognitive benefits arising from the perceptual salience of externally constructed information structures, but also social benefits arising from the interpersonal interactions that accompany face-to-face token exchange.

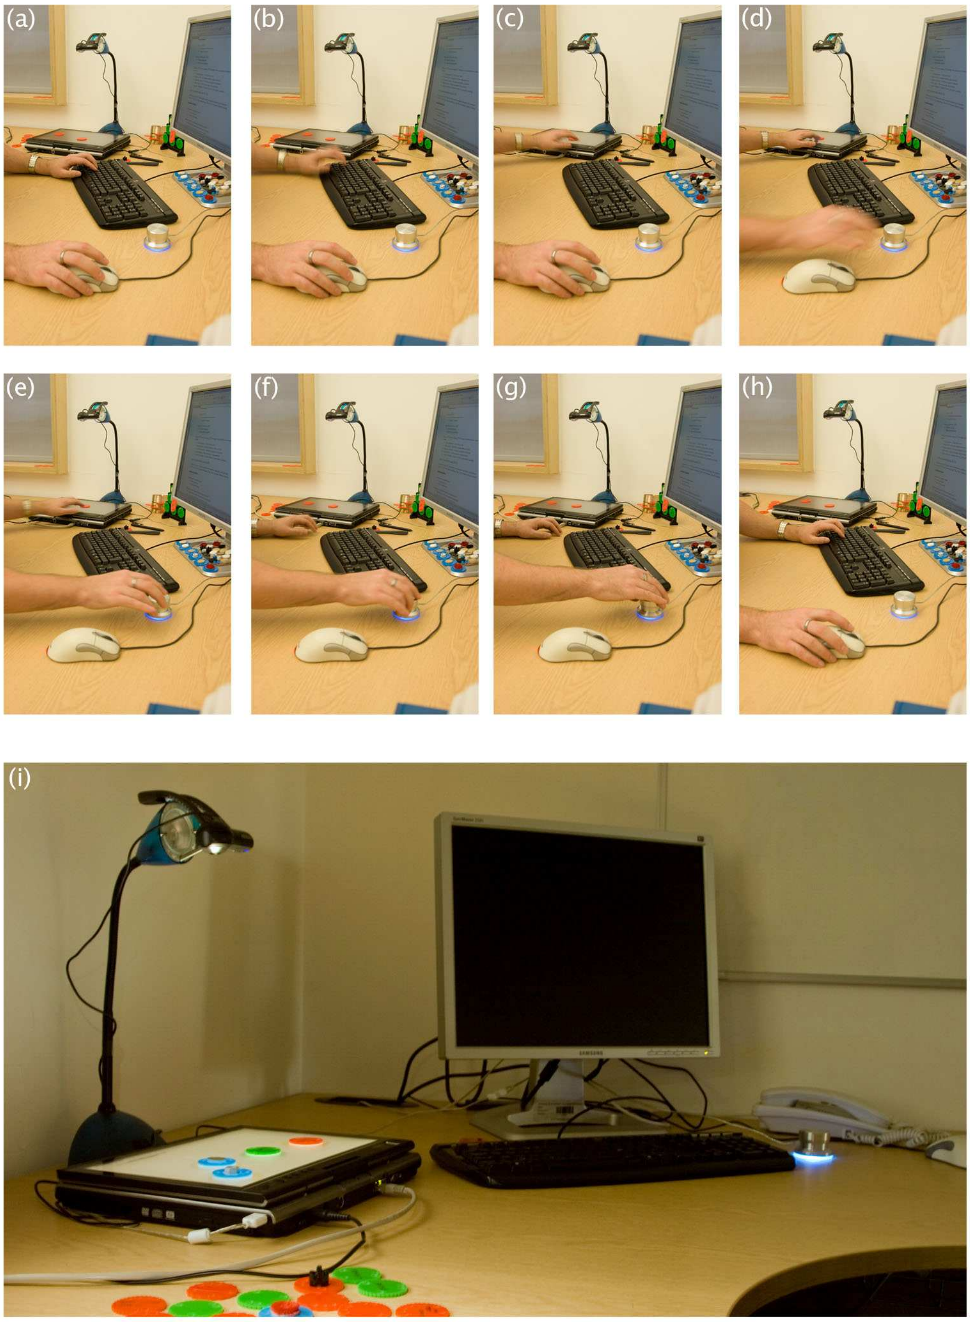

Figure 5.3: Peripheral Interaction Sequence and Interface Structure: (a) both hands in 'home' positions; (b,c) left hand lifts off keyboard and approaches token; (d) left hand nudges token as right hand lifts off the mouse; (e) right hand acquires knob as left hand retreats; (f,g) right hand rotates knob; (h) return to home position; (i) the TUI structure complements the existing monitor, mouse, and keyboard and supports peripheral interactions of the kind shown in (a-h).

Figure 5.4: Appropriation of the Desktop Environment: Users can project their own systems of meaning onto the desktop environment through the systematic arrangement of tokens.

5.1 Overview of TUI Implementation

The prototype TUI implemented has three distinct but related components. The first is the system architecture – how the interactive surfaces and workstations interact with one another and with the server. The second is the hardware selection – how the system architecture is implemented using off-the-shelf physical devices. The third is the token collection – the physical tokens that were iteratively developed until an appropriate form factor and coding scheme was found.

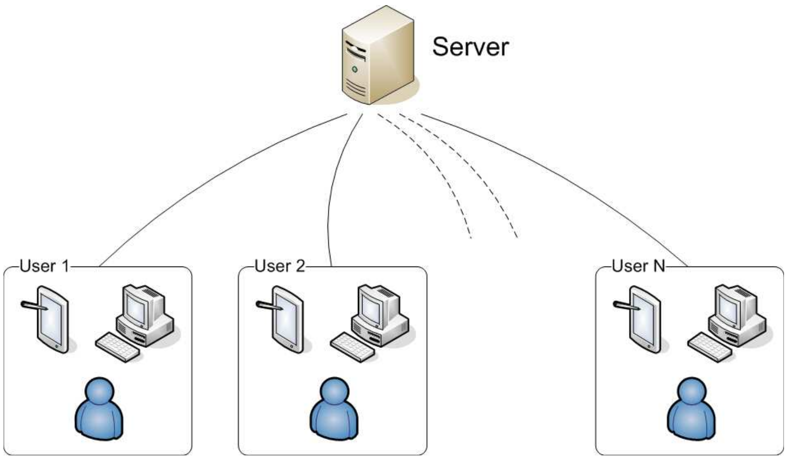

Figure 5.5: Client-Server Architecture, Multiple User View: Multiple pairings of workstation and interactive surface communicate with a single server.

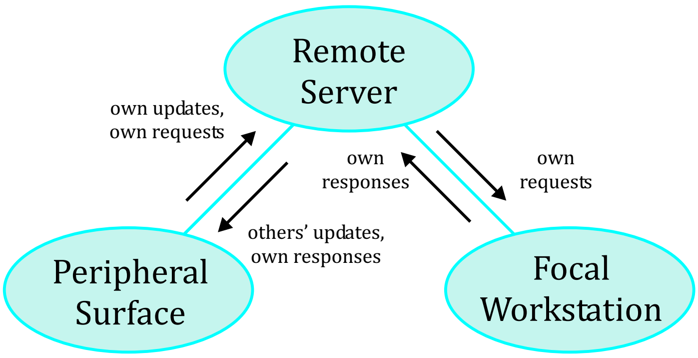

Figure 5.6: Client-Server Architecture, Single User View: The peripheral surface sends requests for user input to the focal workstation, which relays responses back to the interactive surface.

5.1.1 System Architecture

Figure 5.5 shows how the system is distributed in a client-server architecture across three main components: a single remote server, and multiple pairings of personal interactive surfaces and workstation applications. Communication between these components is shown in Figure 5.6. When a user interacts with their personal interactive surface, any updates to their digital state are communicated to the remote server. Similarly, any requests for non-surface input are communicated to the remote server and relayed to the user's existing workstation. This results in a pop-up dialog box appearing in the top-right corner of the user's monitor, in which they can enter text input or use the mouse to select between a variety of options. Figure 5.7 shows screenshots of the user's monitor as they are entering the new task name – the dialog box also contains a button to "unlink" the task token such that it may be 'recycled' and set up as a new task. In a symmetric manner, the response is then communicated back to the remote server and relayed onto the user's peripheral interactive surface. This relay protocol is also used to communicate changes made on the one user's interactive surface to the interactive surfaces of all other interested parties.

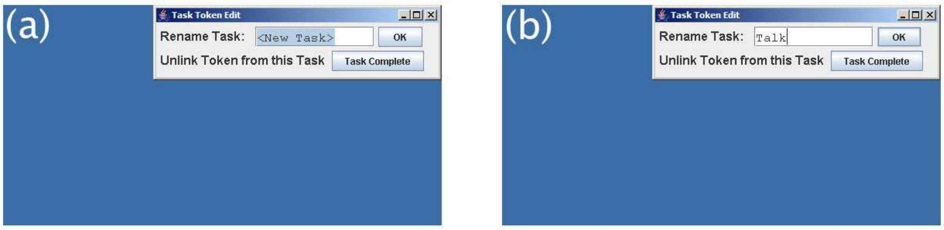

Figure 5.7: Setting a Task's Name using the Workstation: (a) after performing a nudge-click action on a new task token, a pop-up dialog box appears in the top-right corner of the user's monitor; (b) the user types in the name of the task and clicks "OK".

A more detailed description of the system architecture is presented in Section 5.3.

5.1.2 Hardware Selection

The literature review of Chapter 2 presented systems with a variety of sensing mechanisms; for systems of multiple objects, these generally fall into the categories of computer-vision and electromagnetic induction. For the construction of the personal, peripheral, interactive surfaces required by my project, I decided to adopt a computer-vision based approach, elevating a Webcam above a tablet PC and pointing it at the resulting horizontal display surface (Figure 5.3). This configuration makes use of freely available, off-the-shelf components, with a tablet PC conveniently providing a portable display surface as well as computational and networking capabilities. For research prototypes, computer vision has the advantage that it permits experimental, iterative cycles of development without requiring physical re-engineering. For my prototype I use Toshiba Tecra M7-102 tablet PCs, each connected via USB to the Logitech Quick Cam Fusion Webcam and a Griffin Powermate control knob. The application software running on each tablet PC is written in Java, as is the background application running on each user's existing workstation. These both communicate over HTTP with the remote server (similarly written in Java), via dedicated servlets that interface with persistent storage of token data in an Apache Xindice XML database.

5.1.3 Token Collection

The tokens themselves are circular discs, laser-cut from sheet Perspex to about the size of a poker chip (35mm diameter $\times$ 4mm thick) – see Figure 5.8. Each has a unique pattern of holes cut around its edge, so that vision algorithms based on fast radial symmetry detection (Loy and Zelinsky, 2003) can be used to identify and track tokens placed on the surface. Compared to the use of fiducial markers, as in the reacTIVision system (Kaltenbrunner and Bencina, 2007), this method of identification has a number of advantages. From a practical perspective, the combination of a brightly-lit background surface and brightly-coloured yet opaque tokens means that any holes made in those tokens will result in highly contrasting circles of bright light shining through – suitable patterns of such holes can therefore act as an identification scheme for a computer vision system. A secondary consequence of such physical holes is that they are meaningful to the user, both as a means of understanding how token identification works[1], and as a physical affordance for hanging tokens from objects in the environment. Finally, and perhaps most importantly, an identification scheme based on the patterns of holes around the outer edges of tokens leaves their central surfaces free for user annotation. Such freedom is essential if users are to meaningfully interact with tokens away from the interactive surface.

Iterative Prototyping

Figure 5.9 shows the iterative development of this token design, both in terms of physical construction and the means of vision-based computer identification. Photograph (a) shows an example TRIP marker, implemented in Java[2] according to the description given in López de Ipiña et al. (2002). I adapted this codebase to support the generation and identification of TRIP-based markers tailored for attachment on small, disc-like physical tokens; a small sample of such markers is shown in photograph (b). Compared to regular TRIP markers, which require two rings in order to determine the marker's orientation, my modifications were to create compact, single-range markers that did not embed orientation information. These markers were used in the development of my first prototype of an 'interactive' surface: a Webcam elevated above a horizontal, plain white background on which paper-based markers could be manipulated. The outputs and intermediary results from the

[1] Our Tangible Correlate of $\texttt{hidden augmentation}^{[\texttt{hidden dependencies}]}_{<\texttt{TC}>}$

[2] This Java implementation of TRIP was courtesy of Phil Tuddenham, who kindly allowed me to use his code in my initial experimentations with fiducial tracking.

Figure 5.8: Pile of Tokens of All Types: Task tokens (red), Document tokens (blue), and Contact tokens (green).

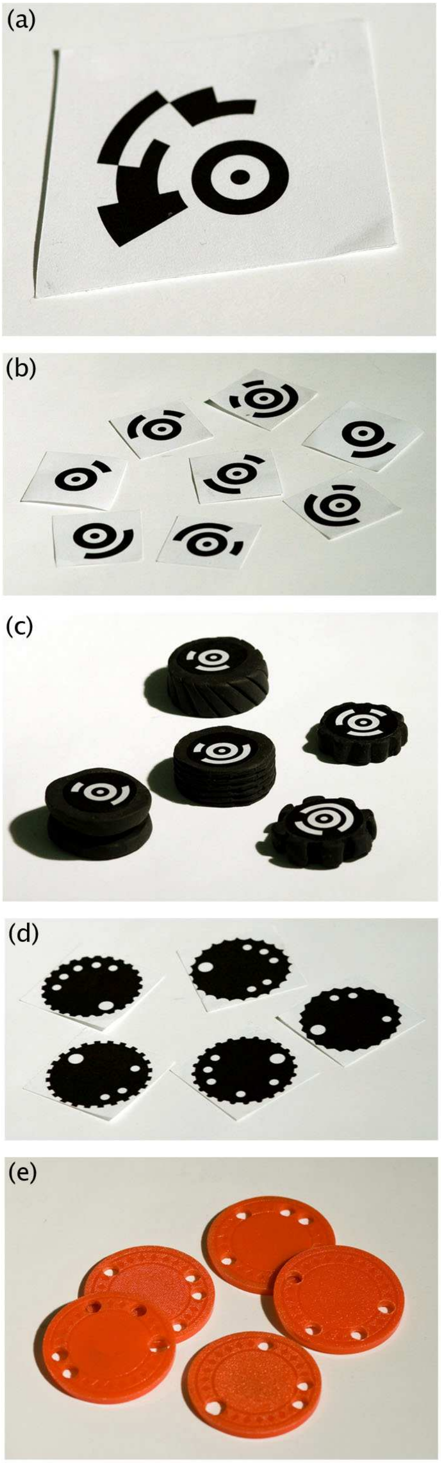

Figure 5.9: Evolution of Token Prototypes: (a) an original TRIP code; (b) my modified TRIP codes – compact, single-ringed circular codes; (c) "3D sketches" using plasticine; (d) paper mock-ups of token codes based on patterns of holes; (e) pre-production plastic prototypes constructed by drilling holes in poker chips.

computer-vision based detection and identification process were displayed on a monitor screen for debugging purposes. This setup was then gradually refined into a true interactive surface, which gave digital feedback in the form of coloured halos projected around and beneath these paper-based markers.

I created many "3D sketches" (Blackwell et al., 2005) of potential token designs in parallel with development of the interactive surface, some of which are shown in photograph (c) of Figure 5.9. During this process of creating and evaluating physical prototypes of tokens, a number of problems with fiducial markers became apparent. These problems – an obscure and technical appearance, lack of space for user annotation, and difficulty decoding in adverse lighting conditions – were all addressed by my coding scheme based on unique patterns of holes, located just inside the circumference of circular tokens. Photograph (d) of Figure 5.9 shows paper-based prototypes of tokens based on such a coding scheme, whilst photograph (e) shows instance of pre-production prototypes, which were created by drilling holes in plastic poker chips. Whilst the hole-drilling approach to distinguishing between different physical tokens has multiple advantages, it also has the drawback that it effectively limits the depth of tokens. For a given interactive-surface size, camera position, token radius, token hole size and number of token holes, the depth of tokens is limited to that which allows an almost undisturbed passage of light from the interactive surface, through a token hole in a corner of the surface, and to the camera. For the constraints of my system this sets a practical limit of around 5mm.

A more detailed description of token detection, identification, and tracking will be given in Section A.6.3. In the next section, I will describe each token type in turn, and how it fits into the broad TUI design.

5.2 Implementation of Token Types

There are three main token types: task tokens, document tokens, and contact tokens. These are supplemented by a special calendar tool used to navigate and control the user's calendar. The design and implementation of both the physical and digital representations of each token type will now be presented.

Figure 5.10: Physical Design of Task Tokens

5.2.1 Task Tokens

Task tokens were dually motivated by ethnographic studies of Air Traffic Control and by investigative fieldwork (Section 4.1.3). In particular, the use of paper flight-progress strips in air traffic control inspired the use of physical objects to track task progress, whilst fieldwork in an office context suggested particular digital attributes for tasks that could, when accompanied by a peripheral style of interaction, address some of the problems of co-ordinated task work within an office environment.

The physical implementation of task tokens is shown in Figure 5.10. They are cut from red sheet Perspex in sets of 20 per user, with each set of tokens having a distinctive edge texture (Figure 5.11). The edge texture of a token set is symbolically associated with the owner of those tokens, providing both visual and tactile cues as to whose task tokens are in each user's possession. The number of task tokens belonging to each user is deliberately constrained such that they become a scarce resource; owners of task tokens therefore need to decide which tasks are most important to the current work context, instantiating and recycling task tokens accordingly. The physical transfer of task tokens can also act as a proxy for the social delegation of tasks. Whilst this aims to encourage conversation around token exchange, it also enables new kinds of social interaction around playful forms of task delegation, such as leaving your tokens in someone else's desktop environment. All such interactions with task tokens are facilitated by the ability to annotate their upwards facing surface as a human readable means of token identification. The use of dry-wipe markers means that these annotations can be wiped off as desired, giving users the freedom to develop schemes of annotation that reflect the history of task and token interactions. The rear surfaces of task tokens retain the original white covering of the sheet Perspex to indicate their correct orientation on the interactive surface.

Figure 5.11: Task Token Edge Textures: Each edge texture corresponds to a different team-member, indicating task ownership.

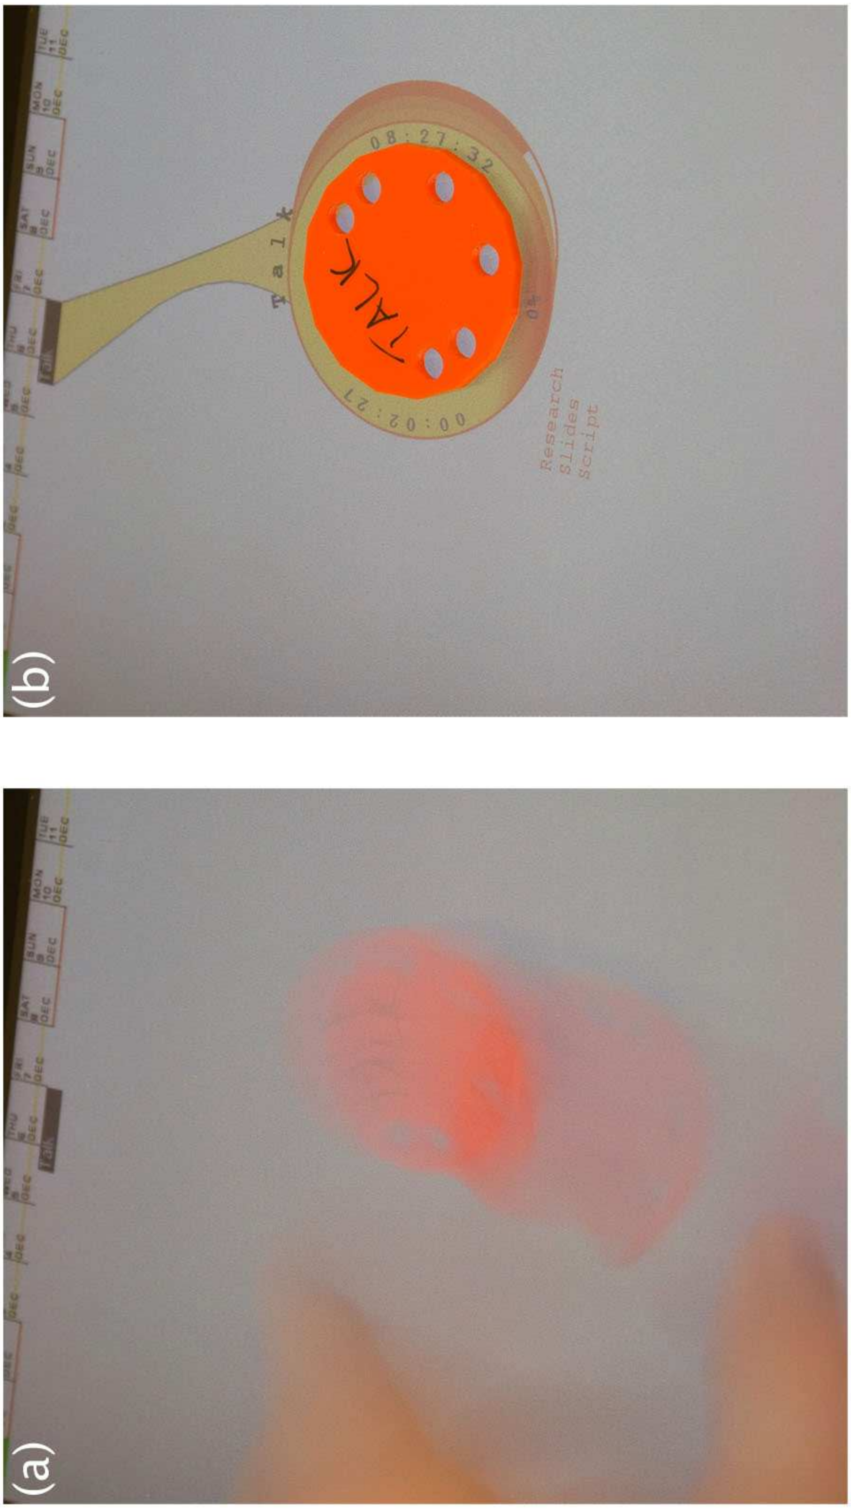

Figure 5.12: Task Token being Tossed onto an Interactive Surface: (a) the task token is roughly thrown onto an interaction surface; (b) a "halo" of information attributes belonging to the corresponding task is displayed underneath and around the token.

Figure 5.12 shows the digital representation of a task being displayed around its associated task token, after the token is tossed onto the interactive surface. When task tokens are placed on the interactive surface, their location and identity are determined by the tablet PC application, and their information attributes are retrieved from the remote server. These attributes are then digitally displayed beneath and around the task tokens, as dynamic "halos" that track the subsequent movement of task tokens around the surface. Task management and time management are closely interrelated, and the digital representation of tasks is similarly coupled with the digital representation of time – a calendar "timeline" that forms the uppermost border of the interactive surface.

The conceptual model underlying the digital representation of tasks is based on three user controllable attributes: *planned completion date*, *estimated work-time remaining*, and an *action list*. For a given task, these attributes are selected by nudging the task token upwards, to the right, and downwards respectively. These rectilinear directions are defined relative to the rectangular interactive surface, with the metaphoric mapping being that to select an attribute of a token, the user must nudge the token in the direction of the attribute as it is displayed on the interactive surface. Control of that attribute is then mapped to the single physical knob (the Powermate) as described in Section 4.2.3.

Figure 5.13 illustrates the initialisation of the task token seen in Figure 5.12. Further elaboration is given in the following sections.

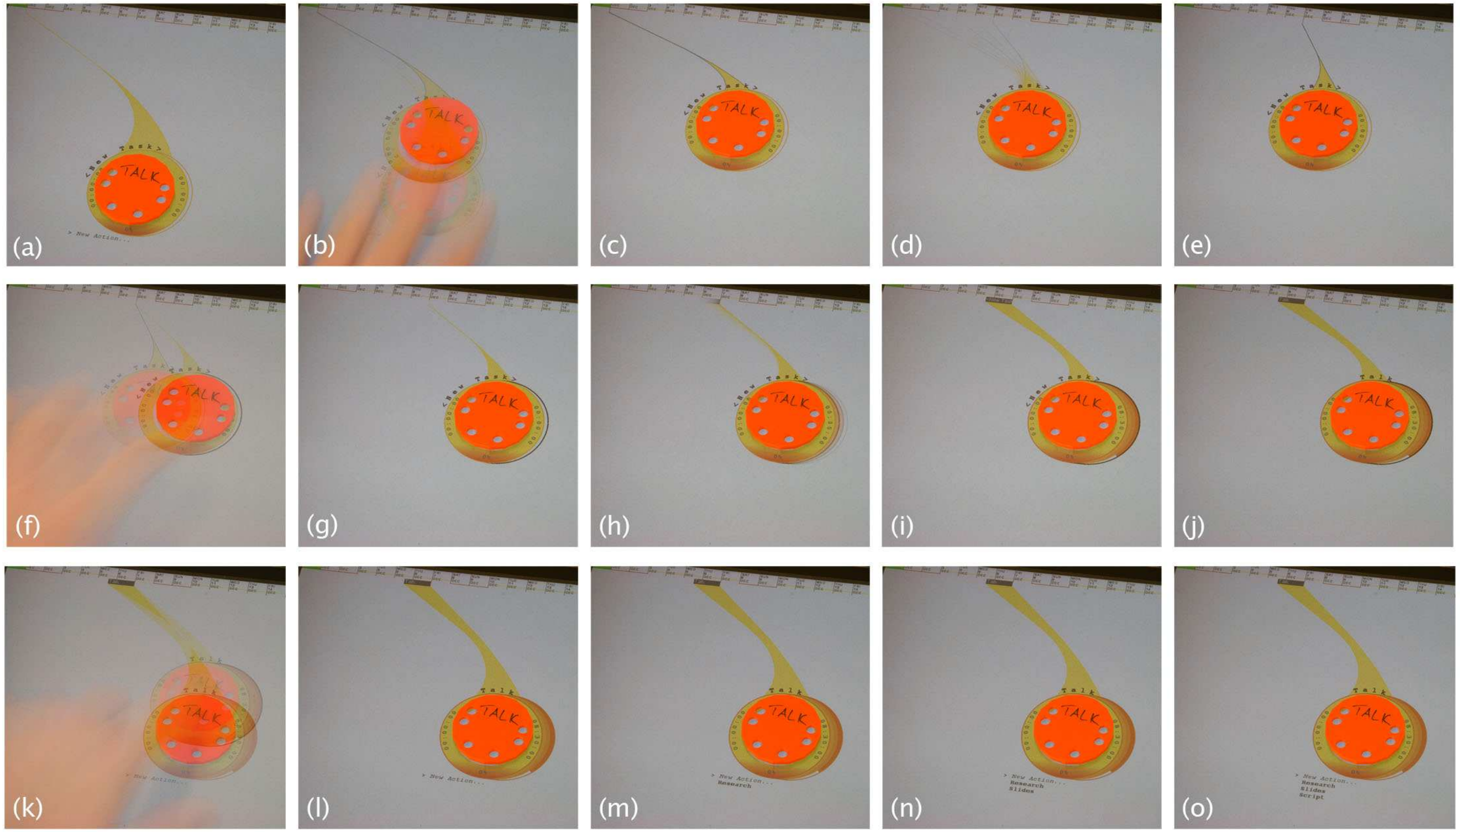

Figure 5.13: Setting a Task's Primary Attributes: (1) nudge token upwards to select its planned completion date (a,b,c); (2) rotate knob clockwise to move the completion date forwards in time (c,d,e); (3) nudge token to the right to select its estimated work-time remaining (e,f,g); (4) rotate knob clockwise to "ramp up" the time remaining (g,h,i); (5) long-click Powermate and type to change name from "<New Task>" to "Talk" as shown in Figure 5.7 (i,j); (6) nudge token downwards to select its action list (j,k,l); (7) short-click Powermate on "New Action..." and type to add actions to this list (m,n,o).

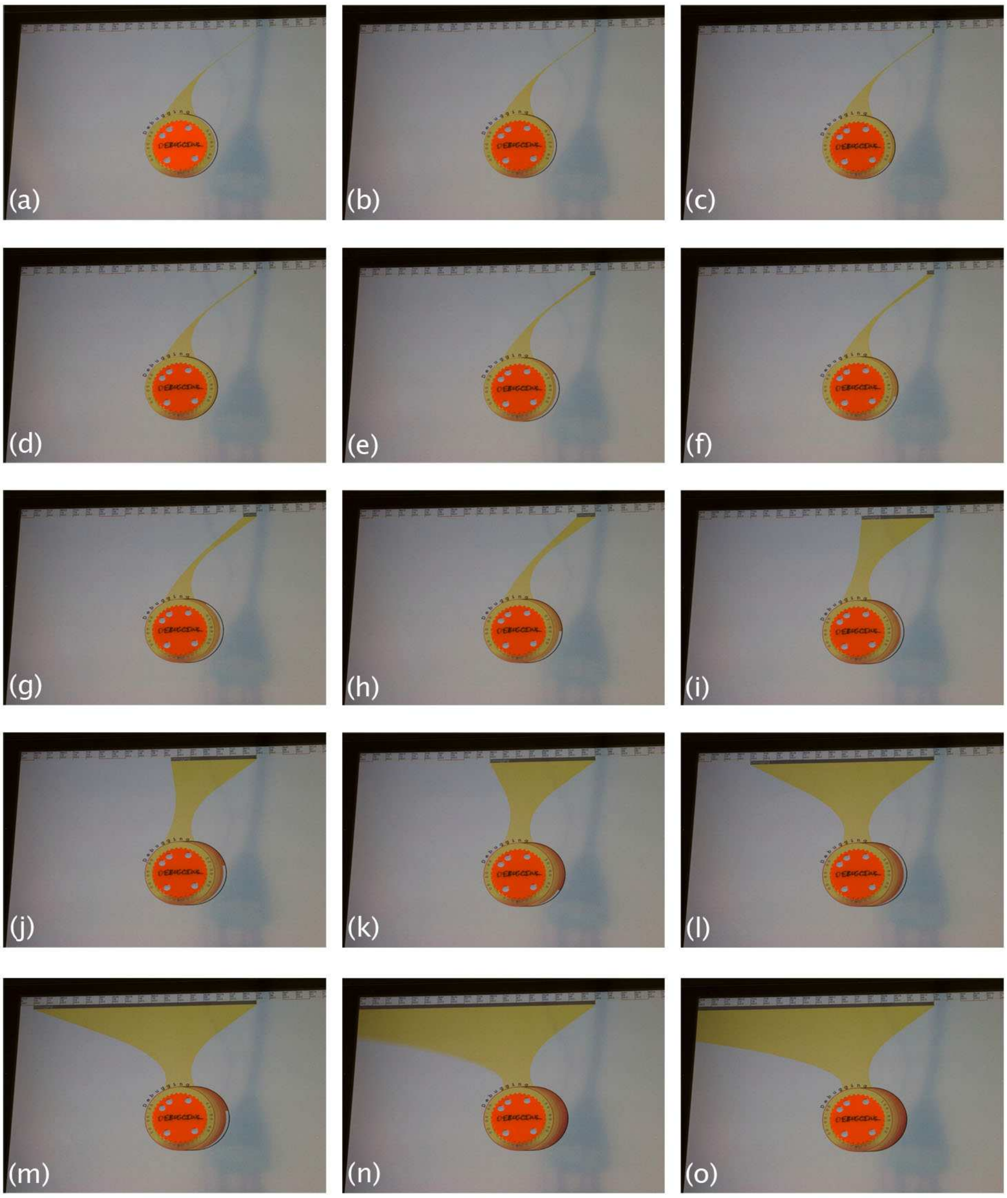

Figure 5.14: Nonlinear Adjustment of Work-time Remaining: (a-o) illustrate the effect of increasing the estimated work-time remaining for a task – the latest restart date of the task is pushed back in time as the semi-circular scales fill and overlap.

Work-time Remaining

The work-time remaining is represented numerically on the right of the halo by a time value, and graphically by a series of five overlapping, semi-circular scales, corresponding to durations up to 1, 4, 10, 40, and 100 hours respectively. Each scale is broken down into 12 increments, allowing time estimates to be specified at a granularity commensurate with their probable accuracy: to the closest 5 minutes, 15 minutes, 30 minutes, 90 minutes, and 5 hours respectively.

Planned Completion Date & Latest Restart Date

Figure 5.14 demonstrates the side-effect that adjusting the hours work remaining of a task has on the representation of its relationship with the timeline. For any task above zero duration, the arc representing the completion date of the task is accompanied by an arc representing the latest restart date of the task. This extends backwards in time from the completion date, and indicates the latest date at which that task should be resumed in order to ensure timely completion. Calculation of this latest restart date takes into account the amount of work-time remaining, the scheduled working hours per day over the course of the task (discussed in Section 5.2.4), and the number of overlapping tasks. The algorithm that performs this calculation is given in Appendix A.3, and assumes an equal division of labour between tasks that overlap in time. Figure 5.15 provides an illustration of this algorithm in action as one task's completion date is moved across the completion dates of two other tasks.

Action List

The action list assumes a model of work in which tasks are conceived as coherent bundles of actions, and where the dependencies between those actions can be remembered by the user, but where the user may wish to remind themselves of certain subset of the task's required actions. The action list attribute therefore allows users to inspect and manage either the next actions to perform, those actions that are currently blocked awaiting responses from others, or any other list of their choosing. In Figure 5.13, photograph (l) shows how nudging a token downwards automatically selects, via the position of the ">" cursor, the pseudo-action labelled "New Action...". A short, downwards click of the Powermate at this point results in a dialog box appearing in the top right-hand corner of the user's workstation monitor. The user enters the name of the new action in this box, which disappears as soon as they press return on their keyboard to communicate the name of this new action back to the interactive surface. The result is shown in photograph (m): the newly created action appears in the action list below "New Action...", with the cursor remaining in position ready to create more actions for the action list. Photographs (n,o) illustrate the addition of two more actions to this action list. Counter-clockwise rotation of the Powermate scrolls the selection cursor down this list, and a short click of the Powermate on any action will cause a dialog box containing the name of the action to appear on the user's workstation monitor ready for editing. Actions are deleted from a token's action list by editing an action name to have zero characters. Rotating the Powermate clockwise will cause the cursor to scroll back up the list, and actions can be moved up or down within the list using press-turn gestures.

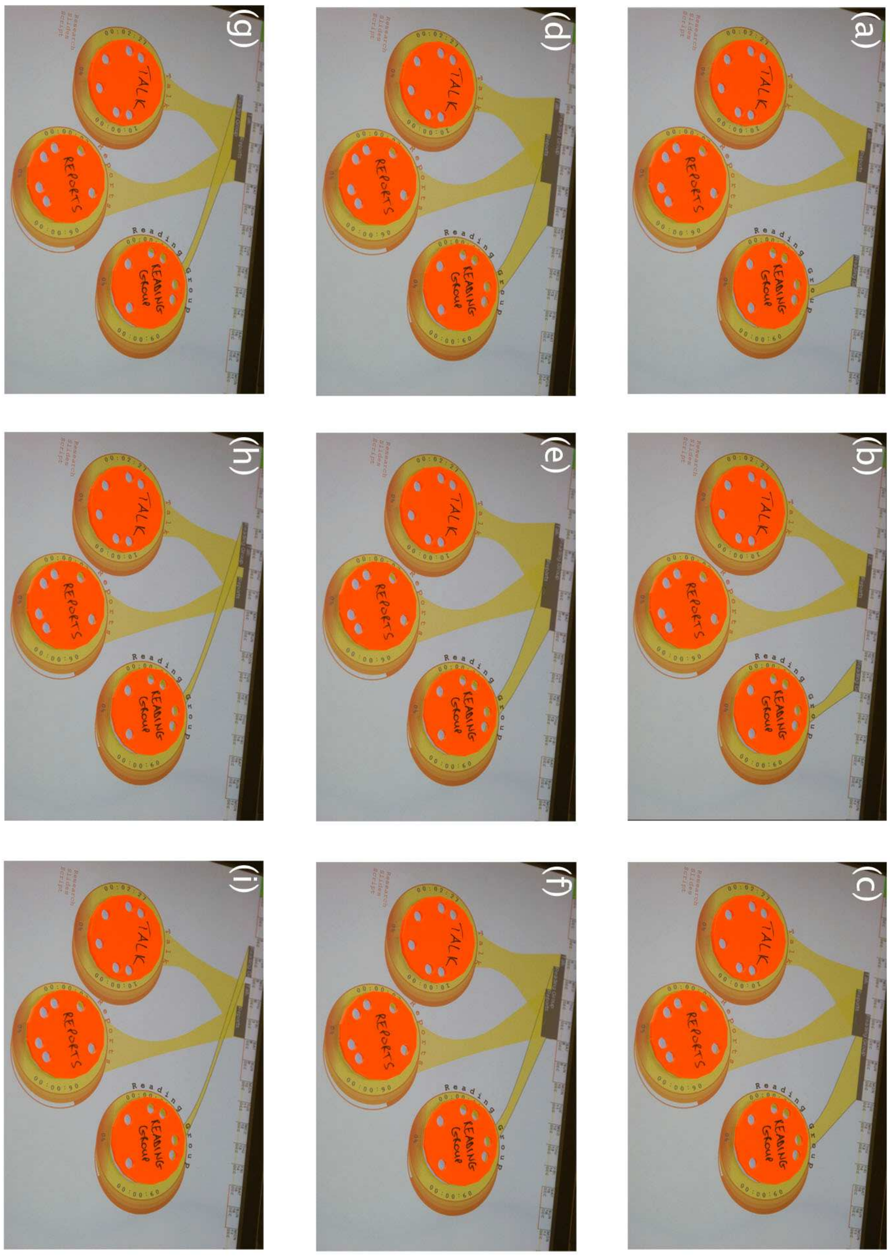

Figure 5.15: Workload Visualisation with Task Tokens: (a,b,c) the planned completion date of a task is moved backwards in time until its latest restart date begins to overlap with the timeline representation of two other tasks; (d,e,f,g,h,i) the latest restart dates of all overlapping tasks vary according to the degree of time-sharing.

Derived Attributes

Whilst the completion date, work-time remaining, and action list of a task are its primary attributes, displayed by its halo and controlled through nudge-turn gestures, the halo of a task also displays secondary attributes that are derived from these primary attributes. The *latest restart date* of a task, already discussed, is one of the two main derived attributes for task tokens. The other is *actual work-time completed*, which is updated in real-time for the task, if any, that is actively being worked on. The fact that a task is active is indicated by the user nudging the desired task token in any direction other than downwards, and performing a short click with the Powermate. This toggles a timer that dynamically counts down from the *estimated work-time remaining*, and increments a counter that accumulates the total time spent so far on the task as the *actual work-time completed*. This value is displayed on the left-hand side of the task token's halo, opposite to the numeric representation of the work-time remaining. To convey the dynamism of this timing process more saliently than by the simple transitions of numeric values, the currently active token is animated by a series of coloured rings that pulse concentrically away from its centre. At most one token can be active at a time, and performing the nudge-click timing gesture on a different token will automatically toggle the timer off for the previously active task token, in addition to toggling the timer on for the newly active task.

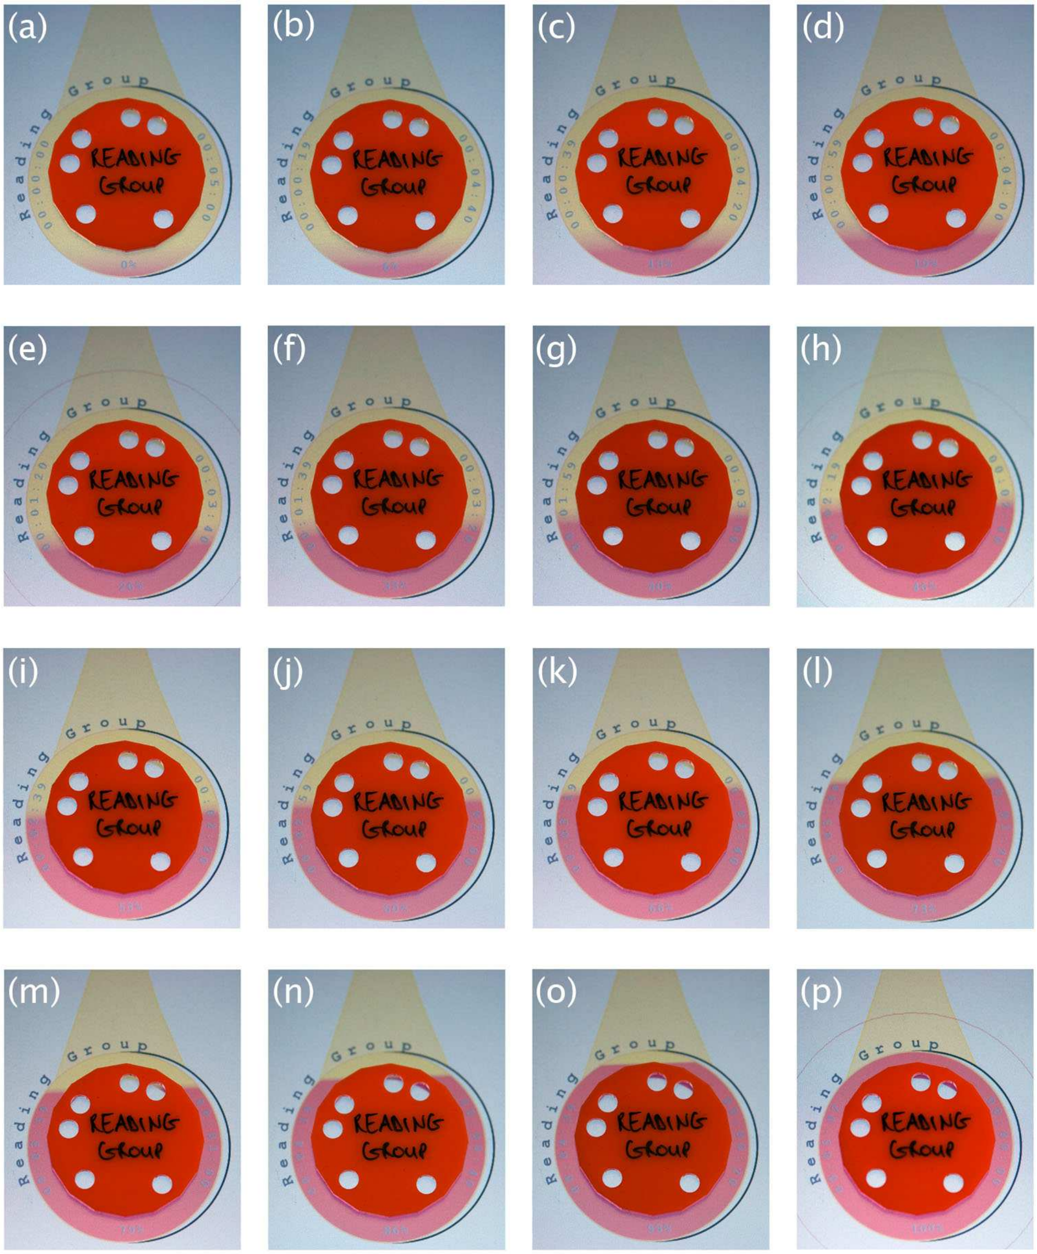

Figure 5.16: Time Tracking with Task Tokens: (a-p) show snapshots at 20s intervals of how the appearance of a task-token halo changes as its "estimated work-time remaining" (shown on the right of the halo) counts down from five minutes to zero. Time accumulates into "actual work-time completed" (shown on the left of the halo). As the derived percentage complete of the task (shown at the base of the halo) increases, the gradient fill of the halo background transitions from yellow to red.

Figure 5.16 illustrates the changing nature of task token halos over the lifetime of the task. At the base of each is a dynamically calculated value of percentage completeness, which is also reflected in the shading of the token halo.

Finally, the digital representation of a task's influence on workload is represented by a grey bar anchoring the task to the timeline. This bar repeats the name of the task, and persists even when the corresponding task token is removed from the interactive surface. This is required in order for users to be reminded of tasks they have scheduled but not retained on the surface, and also to provide a view of tasks in which the influence of multiple tasks is closely juxtaposed in a manner that does not require reference to the task tokens themselves. In the event that a task duration is so short that its name cannot be read from its box representation, then hovering the stylus of the underlying tablet PC above the box temporarily extends the box such that its contents may be read fully. This is the only way in which the system utilises the pen-based functionality of the tablet PC, since the dominant-hand nature of pen-based interaction breaks the bimanual style of peripheral interaction that the system is aimed to support.

5.2.2 Document Tokens

As with task tokens, document tokens were also motivated by both published accounts of ethnographic studies and my own investigative fieldwork. Here, the influential ethnographic studies were not on the use of task-like flight strips air-traffic control, but on the use of paper in office environments. Physical and digital media have complementary sets of advantages, and so the creation of a tangible medium that strongly couples the two could potentially exploit the positive aspects of each media type without necessarily incurring the drawbacks of either.

Figure 5.17: Physical Design of Document Tokens

The physical implementation of document tokens is shown in Figure 5.17. They are cut from blue sheet Perspex and are plain circular discs, with no distinguishing edge textures. They are shared between all users, who take them as desired from a single tray containing both document tokens and document-token materials. These materials are used for rapid recognition and identification of tokens in the physical environment, and attach to tokens via circular recesses cut into the centre of tokens' facing surfaces. A variety of attachment mechanisms are possible, as illustrated in Figure 5.18. One such attachment material – silicone gel – is shown in Figure 5.19. Unlike task tokens, document tokens do not retain the white cover sheet of the Perspex on their rear surface, allowing annotations to be made on the rear surface of document tokens for times when the material mnemonic fails.

Although the simplest form of document token would link to documents on the users' workstations, this would preclude sharing of digital documents through the passing of document tokens. One solution is to only link to documents stored on a shared server. However, my design of document tokens goes beyond the use of tokens as simple handles to documents, and considers the social aspects of document management within workgroups. Rather than linking to any file type, my document tokens link only to Web-based, collaborative text editors. They are not just for document access, but for document access-control. In order to access a particular document, the token linking to that document must be on the user's surface when they attempt to access it. Figure 5.20 shows the creation and access of documents using document tokens, whilst the cloning and sharing of a document using document tokens is shown in Figure 5.21.

Cloned document tokens are equivalent and can all be used on any interactive surfaces to access the same document. This is done by placing the document token on an interactive surface and nudging it in any direction, before performing a short click with the Powermate. This sends a request to the remote server, relayed to the user's main workstation, to open a new tab in the Firefox web browser pointing at the URL corresponding to the token's document. The authorisation process is as follows:

1. On the execution of both the tablet PC application and the workstation application, the user is required to enter login details.
2. After login details submitted from the workstation have been successfully verified by the server, a one-time URL is sent back to the workstation.
3. This triggers an automatic loading of the Firefox browser pointing to the one-time URL, resulting in the receipt of a cookie that serves to identify the user for the duration of the interactive session.
4. Whenever a user adds or removes a document token from their surface, this change is immediately communicated to the server, which maintains the numbers and identities of all document tokens residing on all surfaces.
5. The Web-based collaborative editor functionality is provided by the open source SynchroEdit software. The built-in authorisation protocol of this software has been overridden such that the details of the user currently attempting to access a document is checked against those users that are currently eligible to access that document, by virtue of having his document token on their interactive surface.

Figure 5.18: Document Token Materials and Attachment Mechanisms: The shared materials that can be attached to tokens for identification purposes can be bound to tokens via "Fuzzy-felt" backing, Velcro, white-tac, or through the use of adhesive materials such as silicone.

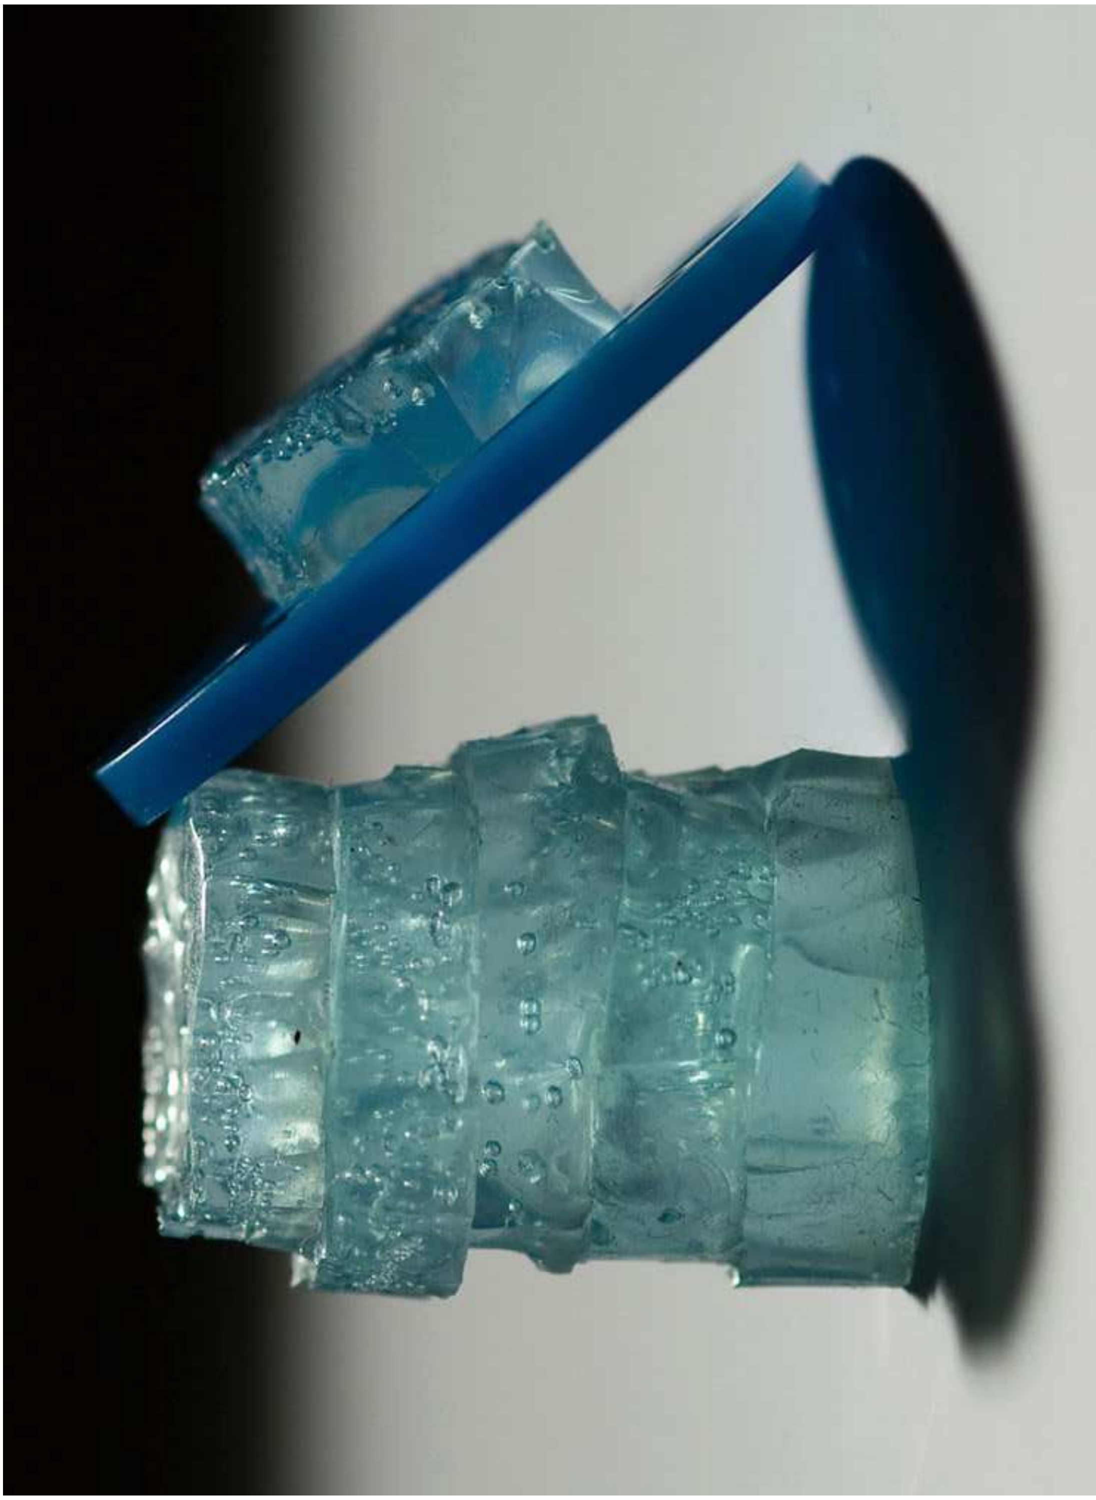

Figure 5.19: Example Attachment Material for Document Tokens: Silicone Gel

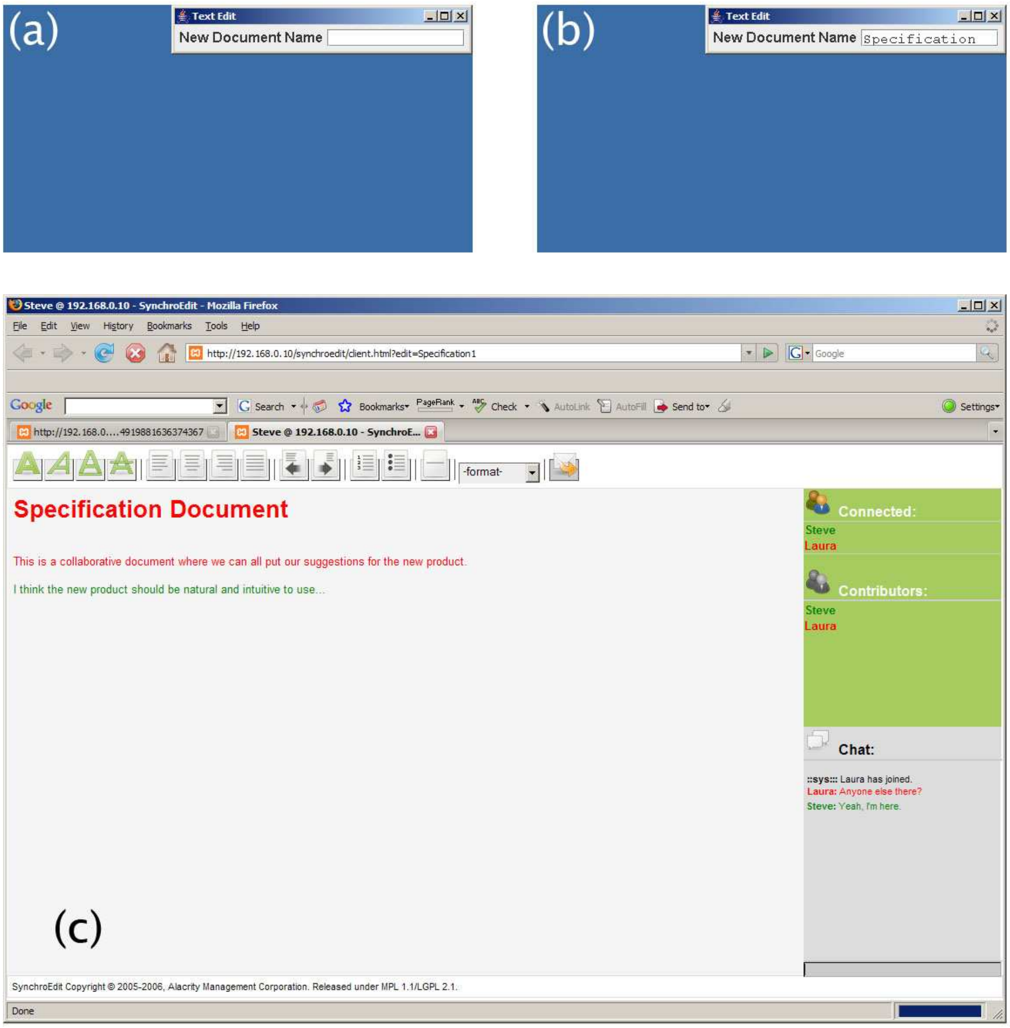

Figure 5.20: Workstation Use in Document Creation: (a) after performing a nudge-click action on an unlinked document token, a pop-up dialog box appears in the top right corner of the user's existing monitor; (b) the user types a name of the new collaborative editor document using their existing keyboard; (c) the document is then created on the remote server and subsequently accessed by performing nudge-click actions on the document token, which causes the document to appear on their existing monitor in a new tab of the Firefox Web-browser. Different text colours indicate authorship within the shared document, and a chat window in the bottom right corner enables real-time communication about the document.

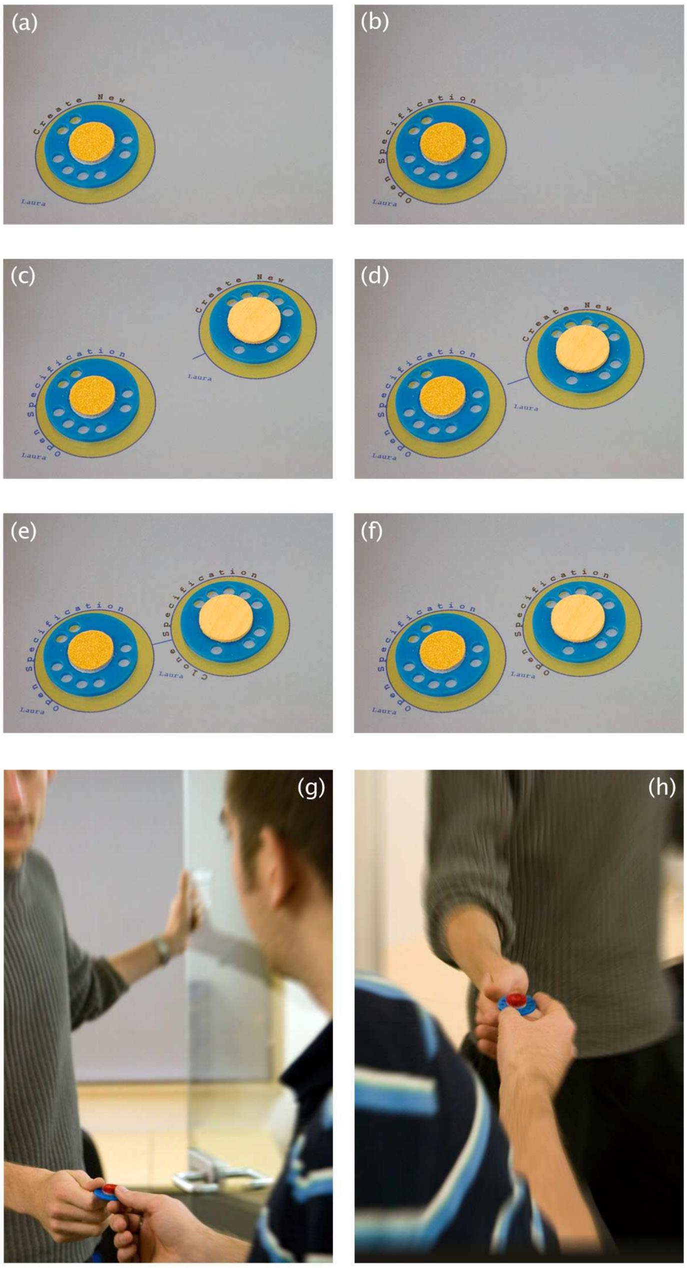

Figure 5.21: Document Cloning and Sharing: (a) an unlinked document token on an interactive surface, labelled "Create New"; (b) the effect of performing a nudge-click action on the token to set up a new document on the server called "Specification" (shown in Figure 5.20); (c) adding a new unlinked document token to the surface; (d) as the unlinked token gets closer to a linked token, a line extends between them; (e) the extended line touches the linked token and its label changes to "Clone <Document Name>"; (f) the linked document token is cloned via a short click with the Powermate; (g,h) document access is shared by passing cloned document tokens.

The rationale behind using tokens as a means of social access control is two-fold. Firstly, the digital representation of document tokens can incorporate a list of all users with the same document token on their surface (this is shown in Figure 5.21, in which the single user "Laura" is currently cloning a token linking to the "Specification" document). Each document token on a user's interactive surface therefore provides a passive communication channel as to which other users are accessing the document, have recently accessed it in the past, or are planning to access it in the near future. This mechanism provides opportunities for ad-hoc collaboration which are otherwise difficult to coordinate: joint authoring efforts traditionally require either prior arrangement or active communication (e.g. a phone call), which are costly in terms of both time and interruption. Peripheral interaction coupled with social access control allows the coordination of joint authoring efforts to proceed in parallel with other work activities.

The second piece of reasoning behind the use of social access control for documents is that it provides face-to-face conversational opportunities and so encourages social interaction. There is no concept of document ownership[3], only of document-token possession. Anyone in possession of a document token may choose to clone it and share access to the document with anyone else.

Security Implications of Token-based Document Access

The use of physical tokens for document access raises the issue of security: what is to prevent the unauthorised access of documents by people using tokens at your desktop when you are elsewhere, what happens if others remove tokens from your desktop for unauthorised access using the TUI of a similarly unaware third party, and what is to stop people unintentionally 'leaking' document access by accidentally handing over tokens to inappropriate recipients?

One simple solution to the first problem is to ensure that all users lock their computers when away from their desks. This would also partially address the second problem, in that 'stolen' document tokens would be of no use without an unlocked system through which their contents can be accessed. This then raises further issues: what about the malicious destruction or disposal of tokens, and what happens if tokens are illegally taken for use by people who also possess a TUI through which token content can be accessed?

Tokens are only ever physical 'keys' to digital information stored elsewhere, and given suitable software, the mappings between physical keys and digital content could always be reassigned. The digital nature of the information and its restricted means of access also provides opportunities for the logging and subscription-based notification of document access events, as well as the setting of security policies that will block access to the digital document even in the event accidental physical handover. The bimanual interface structure could also incorporate some kind of biometric, fingerprint identification into the single physical knob

[3]In general use, there are no distinctions made between the creator of a document and those who share access to it. However, if all those in possession of a document token unlink it from the underlying document, it is possible only for the creator of the document to regain token-based access to that document. This process is described in Section 5.2.3.

as an additional layer of security, or more simply, document access may require the entry of a password before the digital content is displayed.

However, the system I built for contextual deployment is not secure, and it isn't intended to be. Even more so than with tasks and delegation, the tokenisation of documents for access and sharing is primarily motivated by its social implications. The system should be deployed within the circle of trust shared by a small group of users working towards common goals, as a means of making information more publicly available within that circle. Security policies, passwords and so on may be used at a high level to define permissions of particular tokens, people or workgroups, but none of these should detract from the ease of day-to-day "peripheral" interaction with tokens and their content.

5.2.3 Contact Tokens

Contact tokens are tokens that represent other members of the team or work group, existing primarily to support mutual awareness. Whereas document tokens provide a means of passively monitoring document interest, contact tokens allow the user to inspect and passively monitor the work status and work progress of other users. Each user has a green contact token representing themselves and each other user, with the contact token representing a user sharing the same edge texture as the task tokens of that user (Figure 5.22. This enables users to readily identify the owner of delegated task tokens by referring to their collection of suitably annotated contact tokens.

When a contact token is placed on the interactive surface, the resulting digital "halo" displays the name and work status of the associated user (including themselves), with the text wrapping around the token anti-clockwise from the top. To change work status, a user simply nudges their own contact token upwards, towards their status, and perform a short click on the Powermate. They then update the text corresponding to their current status via a dialog box on their existing workstation.

Another effect of placing contact tokens on the interactive surface is that the timeline corresponding to the other user is displayed above the user's own timeline for comparison. This allows users to passively and peripherally monitor the work plans and work progress of one another as they are updated in real-time, which aims to address the reduced level of mutual awareness that can easily occur in the time between formal meetings. An example of how task and contact tokens can work together to help coordinate work efforts is given in Appendix A.4.

Contact Tokens and Document Access

As mentioned in Section 5.2.2, it is possible for the creator of a document to create document tokens linking to it, even if no document tokens remain linked to that document. When a user places their own contact token on their interactive surface, a list of all previously created documents is displayed at the base of the token halo. This list is selected by nudging

Figure 5.22: Physical Design of Contact Tokens

the contact token downwards, and navigated by scrolling a cursor up and down the list using the Powermate. The contact token can be seen as temporarily acting as a document token that links to the currently selected document. The process for cloning this document token proxy is the same as for cloning a regular document token, as was illustrated in Figure 5.21.

5.2.4 Calendar Tool

Each user also has a special token – a *calendar tool* – that they can use to interact with the timeline. Whereas tokens represent items of interest outside of the system, the calendar tool is mainly an interaction device with unlikely cognitive or social function. Its physical form is a cross between task and document tokens – red in colour with a smooth edge texture and central recess – in order to help it stand out from the larger numbers of these tokens. The functions of the calendar tool are to adjust the number of anticipated working hours (either of all working days or of one day in particular), to adjust the scale at which the calendar-based timeline is displayed, and to navigate backwards and forwards in time by scrolling the visible segment of the timeline. Use of the calendar tool is illustrated in Appendix A.5.

Figure 5.23 shows the calendar tool and a variety of tokens resting on an interactive surface, in what may be seen as a typical (if perhaps slightly 'busy') arrangement of tokens. The amount of visual clutter has been reduced by nudging the "Reports" task token to the left – this has the effect of 'minimising' its digital representation such that only the name of the task remains visible. This effect is based on the metaphor of pushing tokens away to reduce their visual clutter, but keeping them on the interactive surface such that they remain available for attribute inspection and manipulation at a moment's notice.

5.3 Technical Implementation of the TUI Prototype

The previous sections presented the TUI prototype from a user's perspective – what the various token types are for, and how to interact with them to achieve particular goals. In this section, I describe the way in which different aspects of system functionality were achieved: firstly from the general design of the system architecture, and secondly from the specific implementation of a token identification scheme and corresponding identification algorithm.

5.3.1 Software Architecture

Figure 5.24 shows the primary components that exist in the system. I will now describe each of these components in terms of the independent dataflow chains in which they participate.

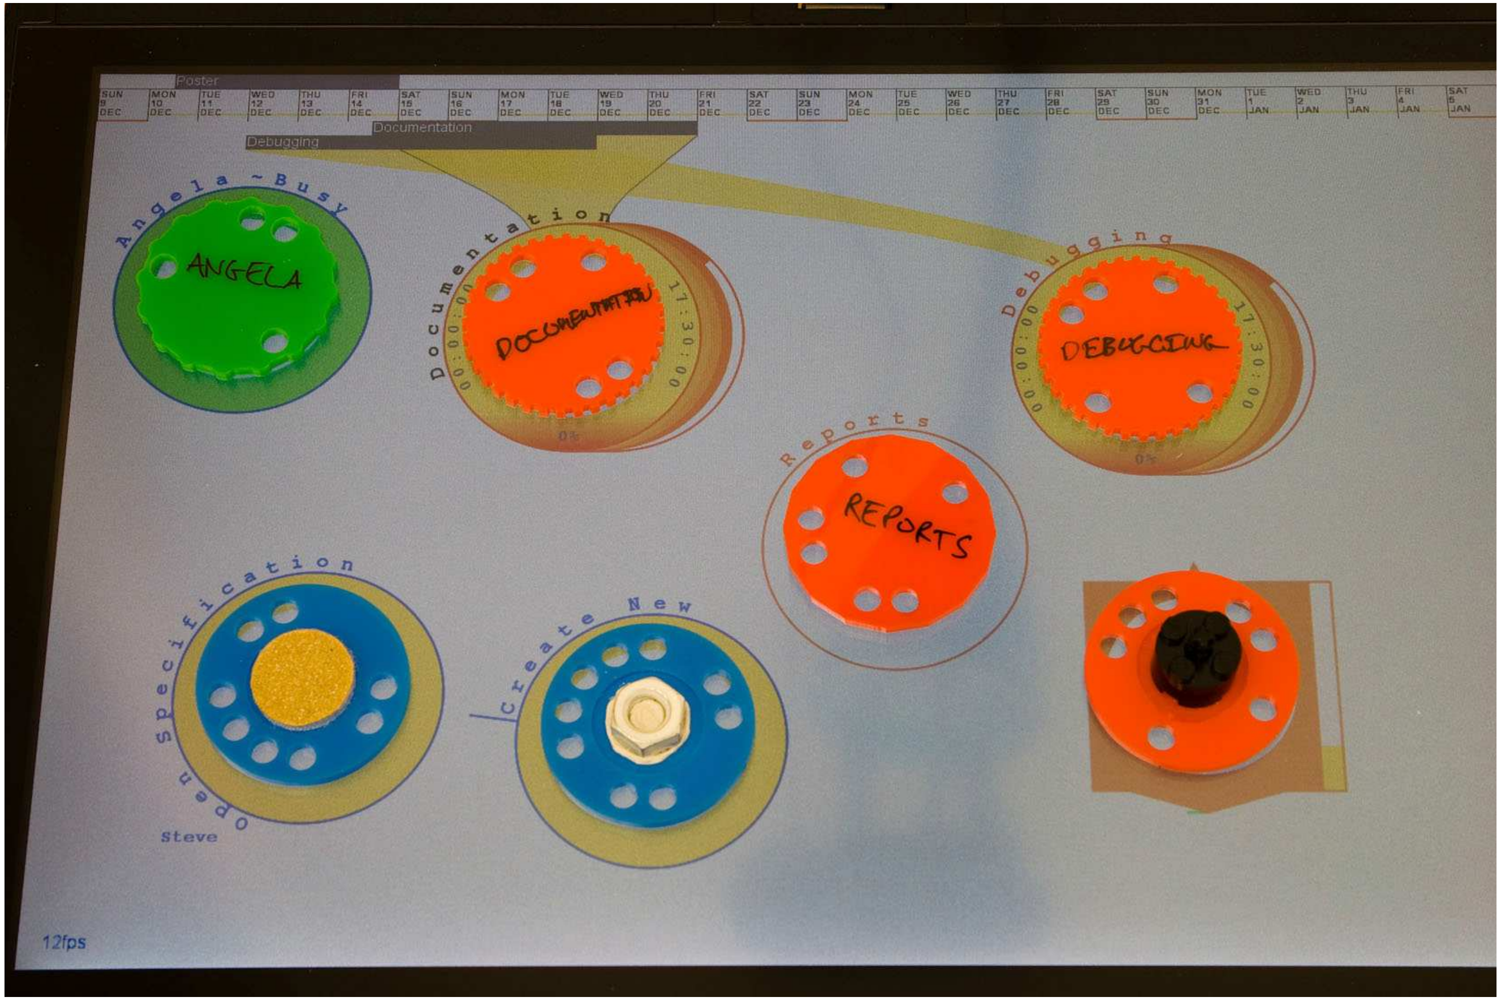

Figure 5.23: Photograph of an Interactive Surface: A typical layout of tokens of all types: contact tokens (green), task tokens (red with edge texture), document tokens (blue with material attachment), and calendar tool (red with material attachment).

Tablet–Workstation Coupling

Initialisation of the TUI as a whole requires that the user's tablet-based interactive surface be associated with their existing workstation, which is used to service input requests requiring either text or an additional layer of indirection for safety purposes, for example when resetting a task token after task completion.

On execution of the TUI applications on both the tablet or workstation, the user is asked to enter their login details in a dialog box. These details are respectively communicated via *Proxies* – called `Server Proxy` on both tablet and workstation – to instances of `Surface Servlet` and `Workstation Servlet`, which send them to an `Authorisation Servlet` for checking. All communication with these Java Servlets takes place via HTTP requests, initiated by the respective `Server Proxy`. Each request results in a new servlet being instantiated by the Tomcat application container to deal with that request, therefore the `Authorisation Servlet` stores user data statically such that it persists across requests.

If the login details are correct, the tablet-based login component will initialise the threads of the `Processing Loop` and the `Poller/Saver Loop`. The login component of the workstation similarly initialises a `Poller Loop`, as well as loading the `Firefox Brower` as described in Section 5.2.2.

Input Handling

The stereotypical Java implementation of the *Observer* design pattern is used to listen for input from the Powermate, with the `Halo Display Manager` propagating press/turn actions to the selected token and token attribute according to a *Chain of Command*. The `Halo Display Manager` acquires a lock on the halo data before changing it, which guarantees these actions will not conflict with other updates to halo data coming from the Server via the `Poller/Saver Loop`.

Input also occurs through the keyboard and the Webcam. Whereas keyboard events are listened for in the same manner as Powermate actions[4], and used primarily to adjust parameters of the computer-vision algorithms, image capture with the Webcam is always triggered at a fixed point within the rendering loop. Interfacing with the *Singleton* Webcam is done through the Java Media Framework (JMF).

Rendering Loop

The first action of the rendering loop is to grab the current frame in the buffer of the Webcam. This image is analysed by the `Processing Loop` (described in Appendix A.6), with the

[4]The Powermate operates by simulating keypresses, with each device action being mapped to a different keypress in the Powermate management application. This application also allows adjustment of the rate at which these keypress events are simulated.

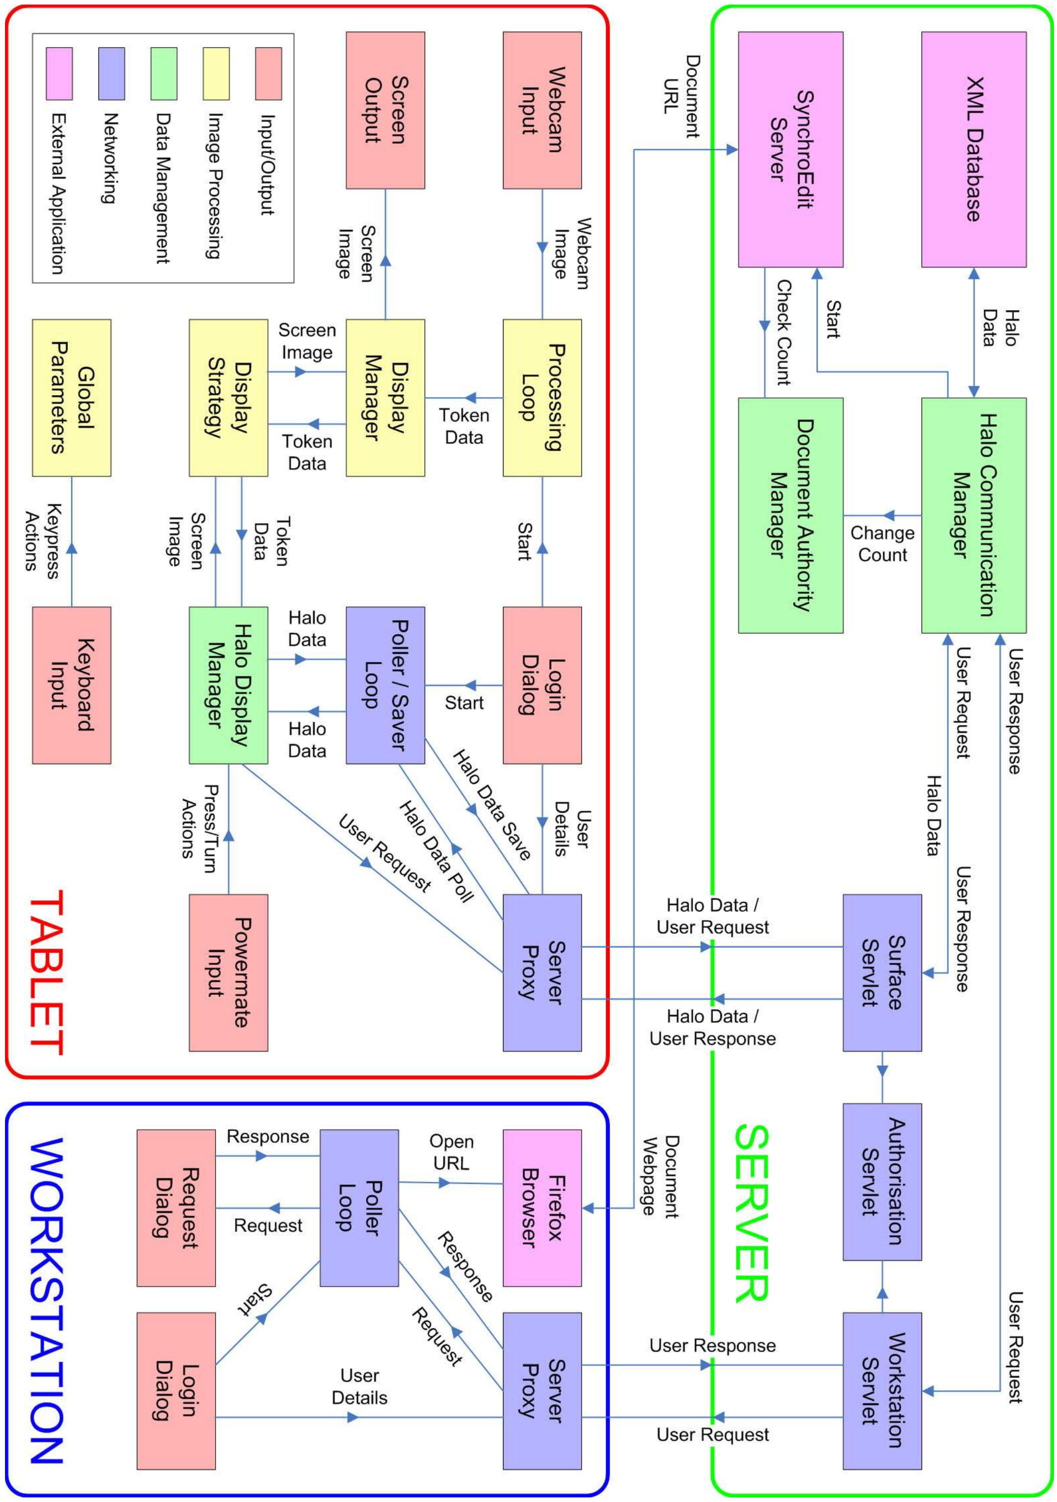

Figure 5.24: Flow of Data Within and Between System Components

resulting token data being passed to the `Display Manager`. This uses a *Strategy* pattern to delegate the display update to a `Display Strategy`, which constructs the next screen image according to the display mode of the system. These are cycled through using the miniature joystick on the casing of the tablet PC, moving from calibration mode, to token-interaction mode, to a variety of debug modes which allow visualisation of token detection and identification (these debug images are used as illustrations of this process in Appendix A.6).

In normal operation – i.e. in token-interaction mode – the corresponding `Display Strategy` delegates screen-image drawing to the `Halo Display Manager`, which maintains the current token selection as well as all other local token data. When asked to create the screen image, it loops over all visible token halos and requests that they return a set of *Command* objects specifying both what to draw, and at what 'layer'. These objects are then sorted by their layer and asked to draw themselves in turn onto the next screen image – this layering ensures the consistent appearance of token halos on the interactive surface, with graphical elements of lower importance being drawn first such that more important elements may be drawn on top of them. An example of this is the yellow swathes linking task tokens to the timeline, which should not obscure the halos of any other tokens, but rather run underneath them[5].

Request Servicing

There are three forms of request that the tablet can make to the server: requests for halo data corresponding to tokens that are unrecognised locally, requests for some action to occur on the user's workstation, and requests to update the halo data on the server to reflect local interactions. These are all relayed from the `Halo Display Manager` on the tablet to the `Halo Communication Manager` on the server, both of which exhibit a *Mediator* role in managing dataflow, and are implemented as *Singletons* for straightforward access to their single instance.

Requests for halo data during rendering by the `Halo Display Manager` are communicated synchronously to the `Halo Communication Manager`, which returns the XML-serialised halo data from an Apache Xindice XML database. In contrast, requests to update halo data held on the server are asynchronously communicated by the `Poller/Saver Loop` to the `Halo Communication Manager`, which saves the updated halo data in the XML database. The `Halo Display Manager` acquires a lock on the database before making changes, which guarantees that user changes are made atomically. The management of which users have which other users' contact tokens on their interactive surface is achieved using an *Observer* pattern.

The final form of request a tablet can make to the server is for some action to occur on the corresponding user's workstation. Requests from tablets and responses from workstations are stored in per-user queues in `Halo Communication Manager`. The head of the request queue is removed and sent to the workstation when the `Poller Loop` polls the server. This request is translated into the appropriate dialog box for user input, with the response being relayed back to the tablet via the `Halo Communication Manager` on the server.

[5]These swathes initially project directly upwards, before curving towards their destination on the timeline. This is intended to reduce interference with other token halos by sweeping around them, rather than tracking underneath them, when tokens are arranged in a horizontal line on the surface. The curves themselves were implemented using four iterations of the Doo-Sabin subdivision scheme.

Document Access

A special case of workstation request is to open the document of a document-token. Whenever a user changes the number of instances of a particular document token on their interactive surface, an asynchronous request is sent to the server to reflect this change in the `Document Authority Manager`. Performing a nudge-click action on a linked document token causes a request to open that document to be relayed via the server to the user's workstation. This request opens a new tab in the `Firefox Browser` pointing to the URL of that document. The Apache Webserver maps this URL to the `SynchroEdit Server`, which checks that the document token count for that combination of user and document is greater than zero before returning the requested Webpage, which contains the embedded collaborative editor document.

Chapter 6

Evaluation by Technology Probe

The previous two chapters have described how a prototype TUI was implemented according to the concept of 'peripheral interaction' developed in Chapter 4. In this chapter, I present details of this TUI's deployment in a real office context, in terms of the underlying methodology adopted, the specific data-collection and analytic methods used, and the results derived from the field evaluation.

6.1 Methodology

The methodology adopted for the evaluation of the prototype was the *technology probe* approach of Hutchinson et al. (2003), which adapts the *cultural probe* approach of Gaver et al. (1999) to focus on the process of co-designing technologies with users. Using this contemporary design methodology, users are exposed to new and provocative technologies in real-world settings, in order to stimulate and focus the ongoing dialogue between users and designers. The intent is not to iteratively hone the probe into a single application or finished product, but to open up new design spaces for further exploration (Boehner et al., 2007).

By their nature, probes are a reflective, interpretive research methodology (ibid.), and as such the influence of the designer and the specificity of the results need to be acknowledged. The implications for design arising from my evaluation are not an objective reduction of the data collected, nor do they represent universal principles that hold across different socio-cultural contexts. They reflect my particular interpretation of users' interactions with the system and responses to it, and serve to highlight what appear to be, in my analysis, the most important considerations when designing a particular kind of interface for a particular class of user, activity and environment. Robson (2002), quoting Sim (1998), notes that this is "sometimes referred to as *analytic* or *theoretical generalization*: 'Here the data gained from a particular study provide theoretical insights which possess a sufficient degree of generality or universality to allow their projection to other contexts or situations"'. It is this kind of generality the technology probe presented here aims to achieve – attempting to characterise a new design space of tangible interfaces for peripheral interaction, based on rich description and analysis of the experiences of a small number of users over many episodes of use.

6.2 Probe Setup

The system was deployed for evaluation at a Cambridge-based startup developing home security systems that integrate with mobile phones, referred to here by the pseudonym "HomeWatch". Three volunteers were recruited following a lunchtime demonstration and talk at their office: P1, the Vice President of Engineering; P2, a senior engineer; and P3, the office manager and personal assistant to one of the company's founders. They agreed to use the system for two weeks in the first instance, with an extension negotiable at the end of two weeks contingent on their experiences with the system in the interim. After two weeks, the volunteers were happy to continue using the system indefinitely, although the start of an intense period of company expansion and physical office reorganisation meant that the study came to a natural end five weeks after the initial deployment.

Although three users is by no means a large sample size for the testing of an interactive technology, the nature of probes is not so much about validating a design as about understanding the design space. Whilst a laboratory exercise would attract a much larger number of participants, the data would reveal little about the real, contextual use of tangibles. As such, the analysis to be presented in this chapter and the next should be seen as just a single step in the process of attempting to understand the costs, benefits, and nuances of tangibility. However, even if the system had been deployed to many more users, I would still have used the methods of data collection and analysis that I will present in this chapter and the next.

6.2.1 Deployment Strategy

Boehner et al. (2007) criticise the literature on probes for its "lack of detail in describing how probes were introduced to participants". The approach I took when introducing this probe was based on the distinction Desanctis and Poole (1991) make between two aspects of interactive systems. Firstly, they define the *structural features* of a system as the "specific types of rules and resources, or capabilities, offered by the system", which "govern exactly how information can be gathered, manipulated, and otherwise managed by users". Secondly, they define the *spirit* of a system as the "general intent with regard to values and goals underlying a given set of structural features"; that is to say, "the 'official line' which the technology presents to people regarding how to act when using the system, how to interpret its features, and how to fill in gaps in procedure which are not explicitly specified" (ibid.).

Such gaps, between the features of a system as they stand by themselves, and the design intentions motivating those features, are traditionally filled by documentation and training. However, given that part of the evaluation process was to observe how users adopted and adapted system features in their own way, care was taken to not reveal the intentions and hypotheses underlying the system during its initial presentation. The tutorial given to participants was therefore a structured, interactive demonstration of the system features only, with no discussion of potential cognitive or social appropriations. This was supported by each user receiving a "complete guide" to the system (Appendix A.7) – printed on the two sides of a single sheet of A4 paper and laminated – containing explanations of the different token types, a table presenting all possible interactions with tokens, and a reference screen shot of the tablet surface. In addition, each user was equipped with a set of 20 task tokens of a distinctive edge texture, a contact token for themselves and the other two users, and a pen with which to annotate tokens. Users' attention was also drawn to the shared trays of document tokens, material attachments for document tokens, and Technic-Lego. The materials were presented as "things you can attach to the recesses of document tokens", and the Lego as "something you can use to manage tokens", but mentions of design intentions such as personalisation, recognition, and play were deliberately omitted in order to avoid potential respondent bias.

6.2.2 Methods

Operating within the technology probe methodology, a number of methods were employed to gather usage data and elicit responses about the system. According to the original paper by Hutchinson et al. (2003):

> *Technology probes are a particular type of probe that combine the social science goal of collecting information about the use and the users of the technology in a real-world setting, the engineering goal of field-testing the technology, and the design goal of inspiring users and designers to think of new kinds of technology to support their needs and desires.*

Each of these goals was addressed by the following data collection techniques in a mixed-methods approach.

Logging

The tablets and the server independently logged user actions and system behaviour, for the engineering goal of field-testing the performance of the system, and the social science goal of monitoring how users support their work activities by making use of their interfaces. The tablets stored data on token recognition and tracking, calendar usage, and on the timings of attribute selection and manipulation; the server stored data on user logons, the creation, name changing and unlinking of tasks and documents, the management of task action lists, the delegation of tasks, the cloning and sharing of documents, and the Web access of document contents. This instrumentation was used to collect data throughout the field evaluation of the prototype.

Interview

In the second week of use, semi-structured interviews were employed to gauge participant attitudes towards the system, to talk about their experiences so far, and to reflect on potential future usage scenarios. These interviews lasted between one hour and one and a half hours, and covered the structure of the interface, the style of interaction, the task, document and contact applications, the use of tokens away from the interactive surface on the desktop and with colleagues, and the impact on individual work, team work and information. The interviews concluded with a discussion of the benefits and drawbacks of the interface as deployed, as well as potential obstacles to adoption for future interfaces based upon it. These interviews were used to address the social science goal of understanding how users thought the interface compared with their previous approaches to managing auxiliary work activities, as well as the design goal of understanding how to build similar interfaces that better address the needs and wants of users.

Elicitation

At the end of the study, as a means of deriving insight into research questions of tangible interface design, structured interviews using a method of card ranking and rating were conducted with each user to gather feedback on the factors they felt contributed most to the overall benefits of the interface. The results were analysed and used to create research diagrams that were later presented to the participants for a final, "graphic elicitation" (Crilly et al., 2006) feedback opportunity. Such presentation of interpretive outcomes to the participants involved in a study is known as *member checking*, which Robson (2002) promotes as a "very valuable means of guarding against researcher bias". The description of this method, the motivations underlying it, as well as the results it generated when applied in this field evaluation, are discussed in Chapter 7.

6.3 Probe Results

The results of the HomeWatch technology probe analysis are presented below in three sections: evaluation of application use; evaluation of the interaction style; and evaluation of work impact.

6.3.1 Evaluation of Application Use

Figure 6.1 shows box-plots of token usage at HomeWatch – indicating the minimum, lower quartile, median, upper quartile, and maximum number of unique tokens seen in hourly intervals. The total number of hours' use – defined as hours where at least one token was recognised and moved on the surface – were 121 for P1, 57 for P2, and 27 for P3. The wide range of total usage reflects the difficulties associated with evaluating a prototype technology in a busy working environment. During the deployment, the system was taken offline for a number of days in order to complete on-site bug fixes and reliability improvements. Both P1 and P2 were away from the office for multiple days at a time, and P3's usage was disrupted because she was heavily involved in managing a recruitment drive as well as an

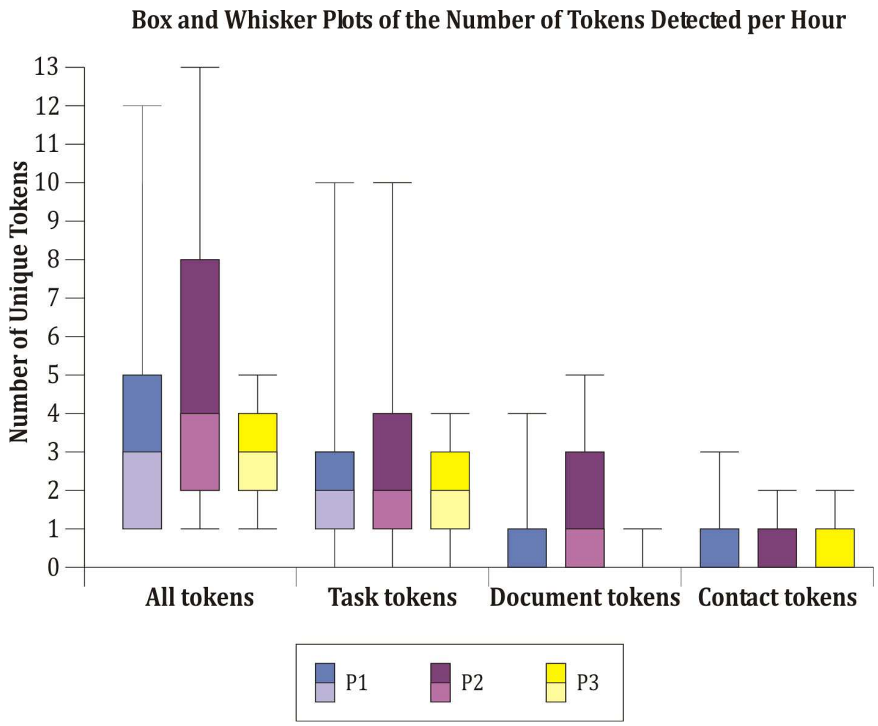

Figure 6.1: HomeWatch Results: Hourly Token Usage

office expansion and redesign. Nevertheless, the five-value summaries given by the box plots provide a reasonable insight into participants' usage of the system.

For all users, task tokens were the most frequently observed tokens. All users also had similar distributions for their contact tokens, since the number of contact tokens was fixed at three. P2 tended to use more task tokens and document tokens per hour, and hence had more tokens on his surface on average. Both P1 and P2 experimented with large numbers of tokens on their interactive surface in any given hour (12 and 13 respectively), but generally preferred fewer (three-quarters of the time, no more than 5 and 8 respectively). This pattern was repeated within the individual token types, with the frequency distributions skewed towards fewer tokens. Whilst tokens can be moved on and off the interactive surfaces during the course of an hour, these values are reasonable approximations to the number of tokens on users' surfaces at any one time, coinciding with casual observations made during the study. Overall, the data show that the system can comfortably work in the region of 1-5 tokens, with spare capacity if necessary. For 'day to day' peripheral interaction the former is likely to be preferable, with more tokens only being used in extraordinary circumstances such as weekly planning sessions.

Tasks

As shown in Figure 6.1, task tokens were the most used token type in terms of the overall time tokens spent on each user's interactive surface. To compare users' interaction patterns in terms of the hour of day in which changes were made, the proportion of each user's total interactions occurring in each hour of the working day were plotted. The results are shown in Figure 6.2, which also breaks down task token usage by feature.

The peripheral management of tasks is the primary interface application, and the chart of task interactions in each hour of the day demonstrates the extent to which task management was conducted in a peripheral manner – i.e. frequent updates throughout the working day. P1 had the most uniform distribution, with peaks at the start of the day, before lunchtime, and at the end of the day. P2 had a large peak spanning two hours at the start of the day, very little activity during the middle of the day, and another peak at the end of the day. For P3 this pattern was reversed – a single hour peak at the start of the day, and a two hour peak at the end. Whilst this doesn't represent near real-time task management updates, 2–3 update periods per day is a significant improvement over conventional weekly updates. Moreover, the fall in usage during the middle of the day corresponds with periods during which users were observed to be predominantly away from their desks – at lunch, or in meetings. The true utilisation of the interface, therefore, is likely to be closer to uniform at the local level – within the morning and afternoon sessions of desk-work – suggesting that users tended to make updates in most of the hours they were at their desks.

The participants' existing approaches to task management were largely self determined due to a lack of imposed organisational constraints, and as such the management model implicit in the digital attributes of tasks was not seamlessly adopted by users, but rather appropriated to match their needs.

The action lists, or to-do lists, were intended to provide a breakdown of the next actions for that task. However, this feature was underutilised relative to task timing, task scheduling, and task estimating (see Figure 6.2). An initial explanation by P1 was that she "couldn't really see them at a glance", but on further reflection she revealed that "the list wasn't really part of the way I was thinking about tasks . . . ticking things off is a bit more hassle than I'd want". On the other hand, P1 noted that "entering a list isn't too bad, and it helps me thinking about and structuring the task", indicating a difference in her willingness to create lists prior to task commencement compared with her willingness to synchronise an action list with the performance of those actions. This is in contrast with P3, who remarked that "it's good to have something that I can either add to or tick off so easily". This distinction is likely to derive from underlying differences in task structure and granularity: whereas P3 had "lots of little things to do", P1 reported having fewer "short tasks" and more "mini-projects" that were ongoing.

The fact that the system was used to support ongoing or repetitive tasks had a knock-on effect on task management in general. P1 thought that the timing was good, being able to "just pull up the graph of accurate time figures that I'd spent on each task, which I could very quickly translate into a timesheet", but would have liked to have been able to specify that a task was recurring – with the same time each week – changing the nature of the task's time attribute from an estimate to a target. Nevertheless, P1 did appreciate the estimate attribute as a way of seeing "how my plans were comparing to reality"; P3 reflected similarly that "it actually forces yourself to think about how you manage your time, and is quite good for how you see your role within the company because you can see what you actually

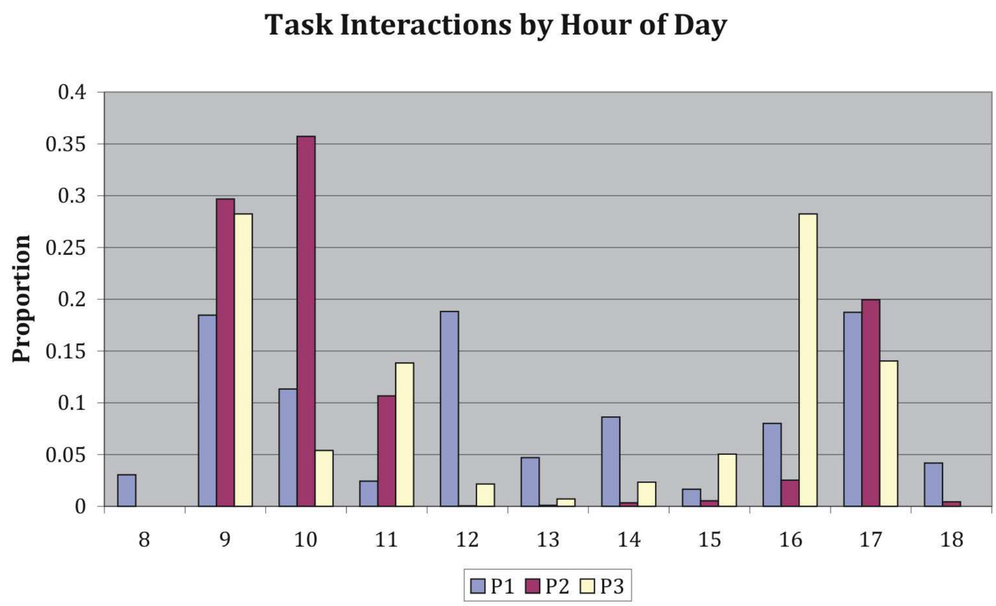

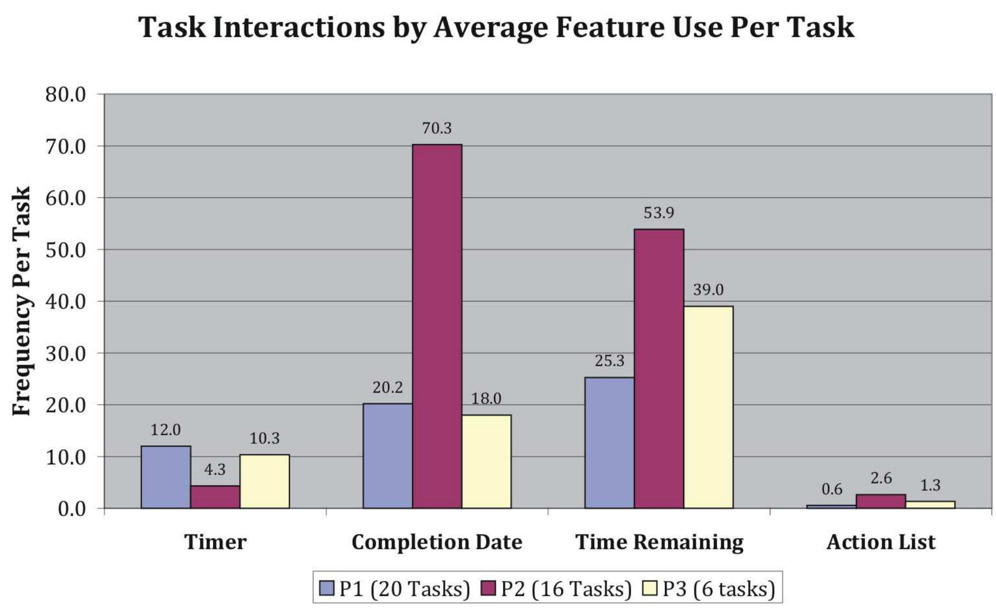

Figure 6.2: HomeWatch Results: Task Token Usage

do". Their positive responses to the task timer were reflected in their average number of timer start/stop toggles per task – over 10 per task for both P1 and P3, representing five or more separate timing periods. The lack of firm deadlines for some tasks also meant that the completion date could be appropriated for purposes other than a firm deadline, with P3 using it to represent "when I'd hope to get things finished by", and P1 using it primarily as a week marker – "at the end of every week reviewing what I'd done and shunting the date across to the next week" – but also as point of reference, used to convey such things as "I said I was going to have this done by Thursday, I should probably pay more attention to this for a while". P2, on the other hand, simply "enjoyed playing" with completion dates, and this was reflected in the fact that he adjusted completion dates three times more often than ether P1 or P3 per task (see Figure 6.2).

The issue of granularity when instantiating tasks as tokens was identified by P2: "setting up a task that is going to take less than an hour probably isn't worth the effort – even up to half a day, it's a bit of a pain in the neck". However, he also supplied the counterargument that "what you think is going to take you half a day often takes you a lot longer", and concluded that "managing tasks in that way, for a start is new for me, but it is useful . . . I've never seem an application that lets you manage your time in such a flexible way before, not even Outlook". As with completion dates, P2 also used the feature to a greater extent than both P1 and P3.

Documents

Figure 6.3 shows a summary of document creation and sharing at HomeWatch. Each document created is represented by a circle, the number inside which refers to the number of sessions that user spent interacting with the document. A session was defined as any number of document accesses not separated by more than five minutes, since users frequently brought an open document into focus through its token, rather than by switching applications to their internet browser, or by switching tabs within their Web-browser, or both. The fact that users often preferred to perform nudge-press interactions, rather than conventional application-task switching, is suggestive of the decreased costs of information access through direct, space-multiplexed tangible interaction.

Whilst P3 created only one document, which she did not share, P1 and P2 both created their own documents as well as sharing four documents between themselves. Collaboration took a number of forms, some of which are presented in Figure 6.4. The document "Design Team Plan" was created by P1, then cloned and passed to P2 as a means of instant collaborative authoring. The following day, P1 and P2 coordinated to edit the same document both in the morning and in the late afternoon, following a day of independent edits by P1. In contrast, for the "Student Notes" documents, the initial collaboration involved the passing of an uncloned document token, followed by independent authoring by P2 before the token was returned, cloned by P1 and then shared by passing the cloned token back to P2. Finally, the second day of collaboration on the "Patent" token involved an initial period of co-authoring, with the token left on each surface for the remainder of the day in anticipation of future collaboration, which failed to transpire. Nevertheless, the passive listing of all users currently

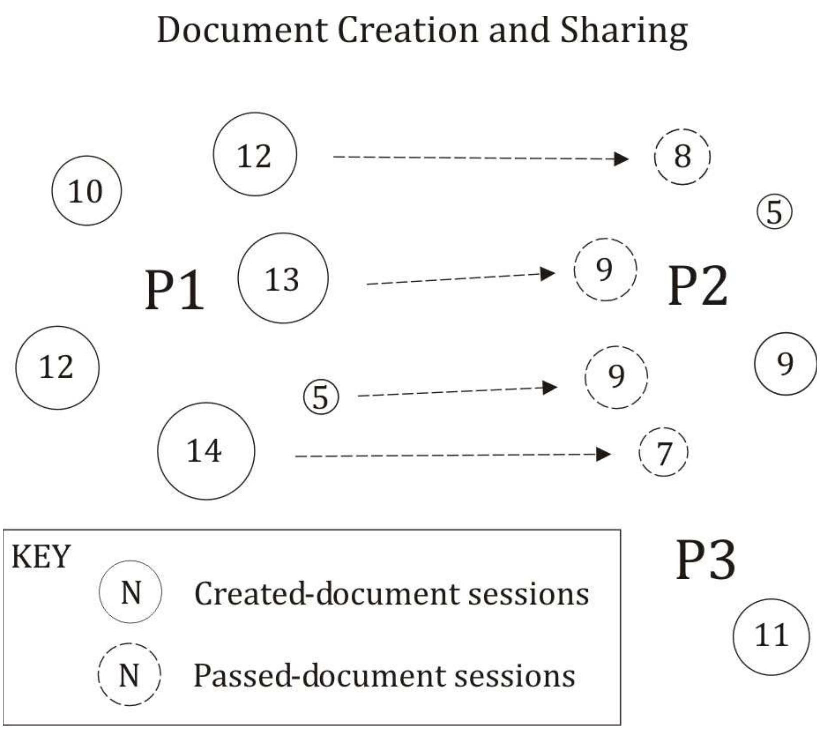

Figure 6.3: HomeWatch Results: Document Creation and Sharing

having the same document token on their surface would have indicated to each user that the other had not taken their token off the surface, acting as a sign that further collaboration could be expected. A design feature that was not implemented, but which would have enhanced this passive communication, is communication via the surface as to when all other holders of that document token last used it to access the document.

In the interviews with users of the system, concurrent access to shared documents was viewed as "a much better idea than version control software" by P1, with the following explanation given by P3: "If you give someone a token, it's just that much more immediate". The immediacy of transferring information through document tokens was echoed by P2, who said that "It's very immediate. You can put it on your tablet, bring it up straight away and spend some time on it, before giving it back. It's actually removing barriers and streamlining the process".

The document tokens introduced an element of "privacy and security" to document management according to P3, whilst also improving accuracy "in the sense that you wouldn't run into problems with people accessing the wrong versions of documents". P1 and P2 both talked about the advantage of being able to physically "hand off responsibility" to one another with regard to who can edit what, using the tokens as a social as well as technological means of access control. Overall, all users saw the document tokens as a valuable way of "collecting and organising ideas" (P1), whether done concurrently, alternately, or privately. The role of local information accessibility was reported as playing a large part in this, with long paths on the file server causing frustration within the company due to the difficulty of

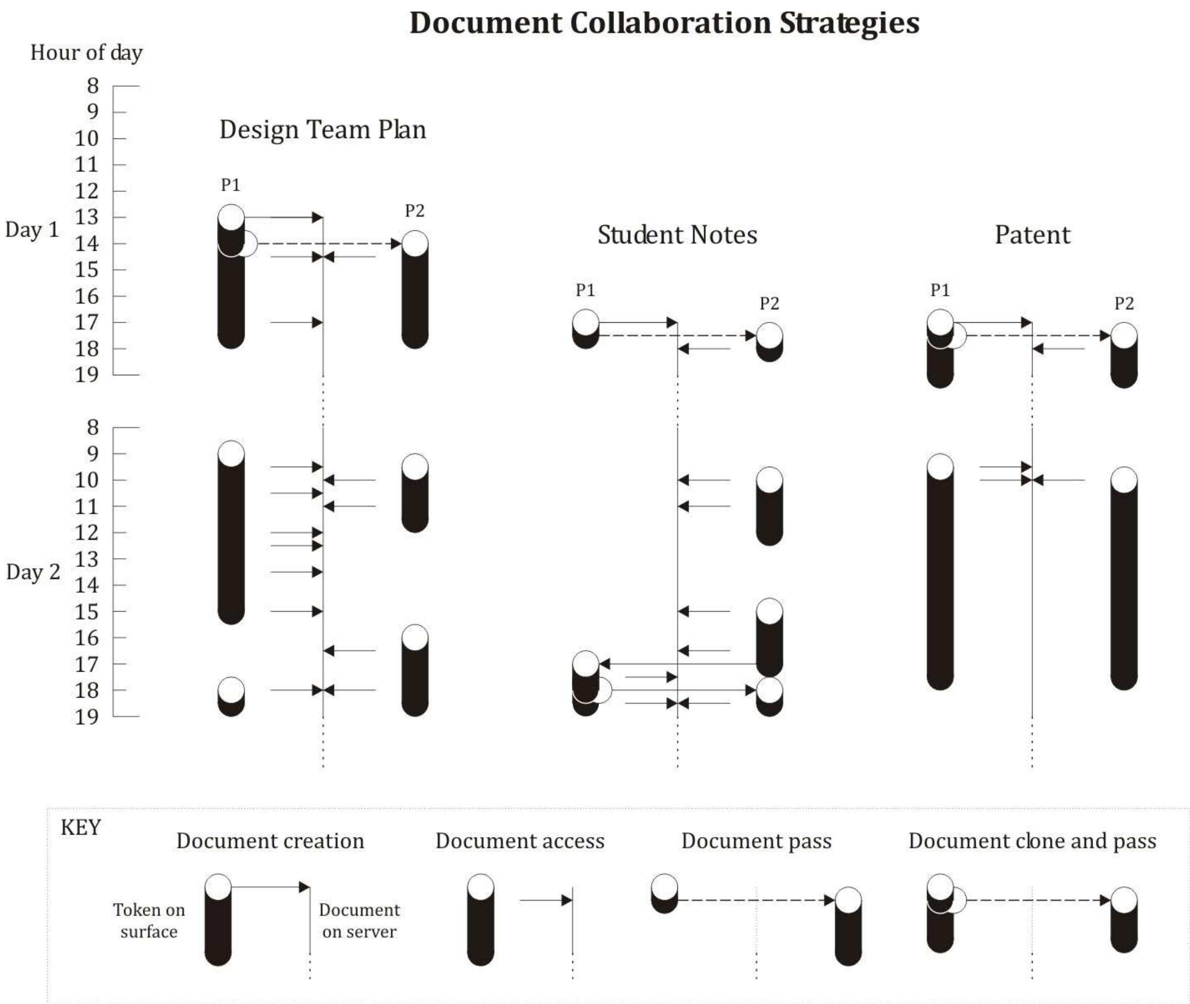

Figure 6.4: HomeWatch Results: Document Cloning and Collaboration

recognising and recalling such paths. The simple extension of using tokens to link to and lock documents on the file server was suggested by P1 as something that would be "really useful". The ability to open regular documents via a file path had been implemented in an early version of the prototype, but was removed after the implementation of collaborative documents since it did not allow for file locking, access control or collaboration. Whilst the ability to integrate tangible tokens with an existing file base would have likely led to much greater use of document tokens, much would have been lost in terms of evaluating the collaborative aspects of document sharing.

Contacts

Contact tokens were used to a lesser extent than the other two main token types, mainly due to the nature of the collaborative relationships between the three participants – P2 reported that he and P1 "constitute a team of two", whilst P3's role as office manager meant that she was "not particularly integrated" with other members of the office. The small but open plan nature of the office also meant that a lack of awareness of one another's work was not the same problem as is often the case in larger, more dispersed teams. Nevertheless, the participants did make some use of contact tokens, with P1 saying that "It's quite fun to see what [P2] is doing. And I like being able to set my status". P3 also identified the value of contact tokens for an office manager or secretary who needs to organise meetings, by being able to see when people are free, what they're currently working on, and what their status is. The ability to set a company specific status message for "direct colleagues" was also regarded as useful by P1, whose status updates were almost exclusively related to the workplace, and included "Tea?", "Meeting", "Nearly conf call time", and "Not building filing cabinets".

The issue of the interface scaling with the number of users was identified by P3: "If everyone had a tablet and everyone had a token for everyone else . . . I guess if you tried to put everyone's token on your tablet, it would get very messy!".

Calendar

Calendar tools were used sparingly by all participants, for a variety of reasons. The digital representation associated with the calendar tool was more abstract than that of task tokens, and it only warranted infrequent use. P1 had to refer to her laminated guide to remind herself how the calendar tool worked, P3 said that she had to refer to the "crib sheet" that was her laminated guide, whilst P2 reported that "some of the lesser used functions, like those on the calendar, aren't quite as easy to pick up".

In terms of usage, P2 logged in at the start of each day, and so each time his calendar view was automatically set such that the current day was always at the extreme left of the screen. He reported that he liked to see "a week or two into the future"; since he could always see this by default, there was no need to scroll the calendar. In contrast, P1 used the calendar tool precisely to counter this automatic effect, to "shift time so that I can see the beginning of the week", and also to "see what I didn't do from last week". All users were happy with the default number of days displayed by the calendar, and whilst users all initially set their own values for "default working hours per day", the fine grained manipulation of individual hours per day was seen as unnecessary in a tool that was for workload visualisation. This feature was seen as belonging to traditional calendars by all users, with P2 saying that "It's all about integrating your application with the user's existing calendar and task list", P3 saying that "If there was a way of linking it into shared calendars, that would be great", and P1 saying that "Provided you were able to tie it in with your regular work calendar, it would really help with widespread adoption". This last point – a marketing observation – points to an inherent difficulty associated with technology probes. The engineering required to construct a system robust and reliable enough to withstand prolonged field deployment can mislead users into treating the interface as a product, rather than a research tool. A competing difficulty is asking users to invest their time in using a system that could potentially cease to work at any point, with no assurances as to whether problems identified would be fixed within any particular time period, if at all. The stance taken with this probe was to make any reliability improvements as necessary, but to leave functional improvements as topics of discussion.

6.3.2 Evaluation of Interaction Style

Whereas the previous section discussed the particular uses of the various token types, this section presents a more general evaluation of the style in which users interacted with their tokens.

Interactive surface

The benefits of the interface exceeded all users' opportunity costs of desk space, with P1 happy to give up "another volume of A4 paper that I can't have a random pile of paper on", and P2 noting that "allocating that amount of space wouldn't be a problem – it's only the size of an A4 piece of paper". The difference between simply allocating desk space and having a tablet PC on your desk was brought up by P2, who thought that the tablet PC had "a bit of a hefty footprint", and that the sides of the laptop impeded the process of sliding tokens on and off the surface. This wasn't seen as a problem by P3, who remarked that "You have to choose what's most important at that point in time, but you can still see enough tasks so that you don't have to always keep moving them on and off the surface", supported by the view that "If you had a massive surface and could just throw everything on there, I don't think you'd be as focused". Both P1 and P2 thought that eight tokens on the surface at any one time was the comfortable upper limit, which would border on being "cluttered", whilst P3 thought that four or five was optimal in order to avoid "confusion". This was reflected by the distributions of token sightings in Figure 6.1. Whilst the interaction capacity of the interactive surface sets a limit on the kind of activities that can be performed on it, the fact that users did not normally operate at this limit suggests that this limit was sufficient to accommodate the general structure of users' work.

Control style

The bimanual control system was well received by all users, with P2 saying that "The nudge to select, the short click to activate, and the long click to edit is actually pretty good . . . I've picked up most of it very quickly and it is very intuitive – it just makes sense, how it works". For P1, "It all seems really smooth", whilst for P3, the simplicity and memorability of interaction was of particular note: "I think it's quite simple, even the interfacing between the tablet and my PC. It's very simple to just move something and just click – that's quite easy to remember!"

The design goal of supporting fast, frequent glances and updates appears to have been met, with P1 looking at the surface "often enough just to refresh my mind about the tokens that are on there", which meant that she was "certainly looking at it about every twenty minutes, probably twiddling with it about every half hour or so". P3 described her interactions as "more a bit of a glance and a quick play around to make sure everything is on track . . . every hour or so".

Switching between tasks also posed no problem for P1 – "I'm putting a token on if I'm doing something different and using my Powermate thing to switch quickly" – nor did eyes-free acquisition of the control knob: "I would just sort of glance over at the tablet to make sure that I was nudging the right token, and the Powermate is just there so I don't really need to look". This was backed up by P2, who said that "It's quite easy to find because it's a big object, and I could see that I'd be able to grab it without looking if I used the system for long enough". This comment by P2 was made after one week of system use, and was followed up by observations of his interaction style. Rather than leading with his left hand to nudge tokens, he was seen to be interacting by acquiring the Powermate *before* nudging, which was contrary to the interaction design. By the end of the second week, however, P2 had increased his confidence in the position of the Powermate to the extent that he was observed to be acquiring it both in the desired eyes-free manner, and in the desired action sequence. The development of such 'muscle memory' after two weeks of use is indicative of the potential efficiency gains for expert users, without compromising the simplicity and learnability of the interaction for novices.

The preceding interview responses to the style of interaction are supported by Figure 6.5, which shows the distributions of users' time differences between nudging a task token to select a new attribute, and using the Powermate to adjust or activate the selected attribute through a turn or press action respectively. The modal time taken (in terms of complete seconds) by users to nudge a token in the direction of the desired attribute, and acquire and operate the Powermate with the other hand, was one second for P1 and P2, and two seconds for P3. This supports the view that the interaction style is well suited to the fast, frequent context switching required for peripheral interaction with digital information. The Keystroke Level Model (Card et al., 1987) can be used as a point of comparison here, with its estimated times of 0.4 seconds to home between the keyboard and mouse, and 1.1 seconds to point with the mouse. Applying the KLM to a hypothetical graphical peripheral interface, the estimated time to acquire and operate the secondary pointing device would be in the region of 1.5s – comparable with the tangible interface selection times shown in the field deployment. Such timings become even more favourable when weighed against the costs

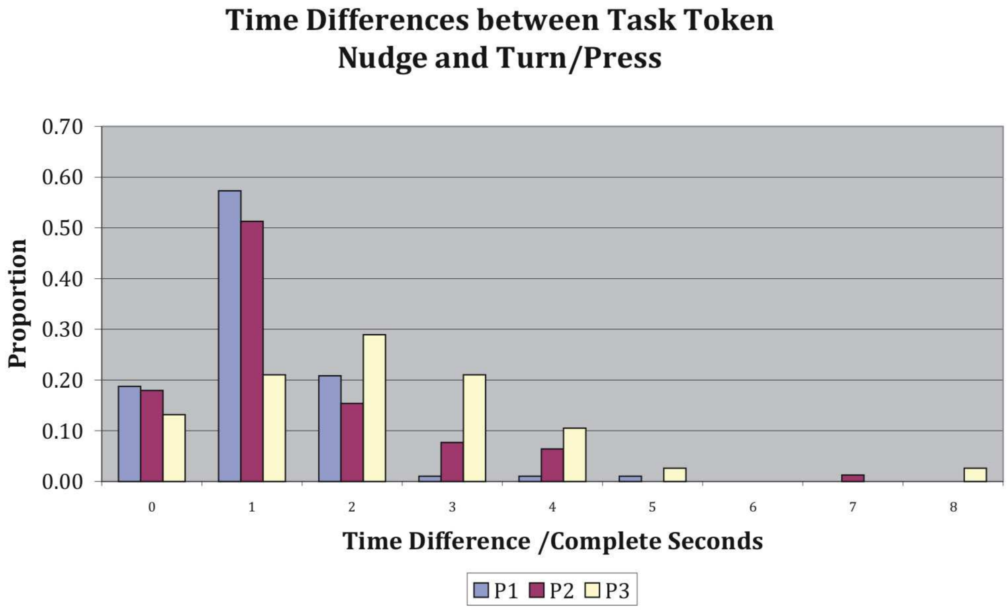

Figure 6.5: HomeWatch Results: Selection/Manipulation Time Intervals

of using a single monitor and mouse to reconfigure a graphical window layout twice – once to switch to a task management application, and again to switch back. In terms of efficiency, using the existing mouse to control a task management application on a peripheral display would probably yield faster target acquisition times, although such an efficiency gain would need to be considered in the context of multiple evaluation metrics, not just efficiency.

6.3.3 Evaluation of Work Impact

Taken together, the TUI structure, the various token types, and the style of interaction have the potential to make a significant impact on the work of both individuals and teams – this section presents an analysis of user comments in these areas.

Individual Work

All participants thought that encouraging them to keep track of their time was the most significant way in which the interface had affected the way they thought about and carried out their work. P1 said that she liked being able to "track things", and while she'd "never been particularly good with my post-it notes", she thought that the "more tangible things . . . really helped" her in becoming more organised. She describes the problem of paper-based lists in the following interview extract:

> The difficulty with post-it notes and other paper lists is that fifty percent of the items on the list get done, and the other fifty percent drag on for ages; then your list looks half done, and you start thinking, "Urgh, it's all manky", and then you think, "Well, I'll make a fresh list – it'll make me feel better", but then you have to duplicate the items over. With individual tokens, you've got less bonding between items, which makes sense because the items are often unrelated. But with lists on post-it notes, the items needn't be connected but they are, because they're stuck together on the same post-it note. And whenever I've tried lots of little post-it notes I've gone mad! With tokens, it's certainly more usable.

P1 also pointed to the fact that sticky notes "just go straight in the bin when they're done", whereas with the tokens "there is a permanent record on the server that I can look at and share". These observations correspond well with my design objective of exploiting the respective advantages of both physical and digital media.

P2 described the model of work-time management underlying the system as a "major shift", with the physical tokens being "just an extra layer on top of that" that assisted the new work style by making it easier to access and manipulate time-related parameters. He also pointed to the trade-off between the "overhead" incurred in setting up a task, and the time savings arising from improved management, concluding that the time investment was in fact "necessary triage [sic]" and that he was "convinced already that it's a boost to productivity". He did, however, make the suggestion that "It would be better if you could create tasks in advance though, all in one go, and then 'pick them up' with tokens as and when you needed them" – essentially reducing aggregate setup costs by batch processing, and suggesting that whilst task management may be conducted in parallel with users' primary work, task creation requires a greater level of concentration.

For P3, the interface had "changed things in the sense that I'm more aware of the time I'm spending on things, which has helped me manage my time better". The benefit of tokens was that she could "put them on the surface and say, 'Right, I'm working on this now"', which was like "committing to doing that work". This was supported by the fact that "you can physically see your time, rather than just a list of things to do", and that "Just seeing it there and having it accessible next to you means you can easily say, 'OK, I've got that to do, that to do, that to do ... I'd better get on it"'. P3 also noted that whilst physical tokens only increase information accessibility within the desktop environment, and could actually decrease the remote accessibility of information, she thought that this "could be a problem, but only a small one", since the nature of the organisation meant that people only needed to think about tasks and documents when at their desks.

Team Work

For P3, the major benefit alongside time management was the increased availability of information: "In terms of the information about what people are working on, how long it's taking them, when it's going to be done and whether they can be disturbed, I think it really helps, because we don't have shared calendars". All participants remarked about the differences between task delegation through token passing, compared with through email or verbal instruction. P1 said that "If I give someone else a token then I don't need to worry about it as much as if I just tell them something, because I don't have to keep asking about it. If I've given them a token, I figure that it should be enough of a reminder for them to do it, and I can leave them alone unless it gets desperately urgent". A rationale for why tokens might make users more likely to complete delegated tasks is given by P2: "When you've got a token, you can deal with it. Making it physical as well as virtual means that you can't just delete a token – well, you can, but it just doesn't feel the same. The sharedness of it all, the fact that other people can see what you're meant to be doing, means that you're much more likely to do it".

This consequence of having tangible representations of tasks – encouraging more timely task completion – also has a "main drawback" according to P2, that "giving your task tokens to others defines a power relationship. You're giving them this concrete thing that they can't just forget about". The reason for the "hierarchical system of delegation" identified by P2 was the ownership of tokens: "If someone gives you a task through one of their tokens, you kind of feel like you're doing their work for them – you're doing the work but they get to claim it". A corollary of this was that P2 felt that "Having tokens that belong to particular people kind of makes it harder to delegate tasks – because a token is your token, giving it to someone else is asking them to do work for you". He suggested a "category of ownerless tokens" or "team tasks that anyone could pick up" to combat this unease, but also noted that the existing system "would work in a formal environment of project management – having your manager's tasks on your desk would probably mean that you spend more time on them". The potential to introduce social tension such as this was also identified by P1: "My boss could also criticise the way I was spending my time, which I wouldn't like very much".

Overall, none of the participants thought that the concept of using physical tokens to represent digital information would pose any problems to widespread adoption. Their only concerns were with the reliability and flexibility of the technology, as described by P2: "If the technology is there and robust, I don't see any problems. Having physical tokens is good in many respects. They do need to be adaptable though, in terms of the different ways people do work. It also needs to be technologically seamless – you can't require the setup of however many bits of equipment, it just needs to be there". These sentiments relate in a number of ways to Weiser and Brown's (1995) notion of calm technology and the disappearing computer – peripheral interaction should support interaction at both the centre and periphery of attention as required, and at all times users should be thinking of the task, rather than the technology.

6.4 Coding of Token Roles

After the fieldwork had been completed, and the interview transcripts analysed in terms of application use, interaction style, and work impact, they were revisited from an alternative perspective – that of the different *roles* that physical tokens can play in peripheral interaction. An understanding and labelling of the different ways in which users can appropriate tokens for cognitive and social purposes could lead to a specialised design vocabulary for describing the goals of any token-based interface.

Six different token roles that were derived from an open-coding of the interview transcripts, by asking the question: What was the user's goal in using tokens in the manner they describe? The coding process itself took the following form:

1. Identify an excerpt from the interview transcripts that identifies a significantly new form of token usage not covered by existing token roles, and tag the excerpt with a provisional code identifying that category of token use.
2. Proceed through the interview transcripts, tagging excerpts with the appropriate usage codes. Whenever a significantly new form of token usage is identified, tag it with a newly created code as in step 1.
3. Once this phase of the coding process is complete, revisit untagged parts of the transcripts attempting to retrospectively apply tags that were created in the first pass.
4. Create a new text file for each provisional code, and copy into each the corresponding transcript excerpts. Copy those excerpts tagged with multiple codes into the multiple corresponding files and mark as duplicates.
5. Examine each file of related excerpts for signs of internal structure, and break into multiple files as appropriate. Give each a new, more refined one-line summary of the token role it represents.
6. As the definitions of token roles are refined and revised, the sharper boundaries between them should mean that most duplicated excerpts become more representative of one token role than the others. Remove less relevant duplicates until each excerpt belongs to a single token role.

The initial coding of token roles identified tokens being used as "interaction devices", as "cognitive aids", and as "social currency". The respective goals of these categories were to interact with digital information, to 'informationally jig' the desktop environment as an aid to memory and cognition, and to provide opportunities for conversation and information exchange.

The generality and granularity of these initial codes resulted in many excerpts from the interview transcripts being duplicated across these codes. One of the main problems was that both digital and social interaction have cognitive components; another is that there are multiple elements of cognition and multiple physical strategies for supporting each. For example, a token can be given a visually salient attachment or put in a visually prominent place to act as a reminder, but the physical form of tokens can also be used to indicate its social meaning, and the location of tokens can also be used to support rapid interaction with digital information.

Token Role	**Description**
Digital Instrument	*To create, inspect, and modify digital information through actions on physical objects.*
Knowledge Handle	*To offload information and uncover action possibilities through manual manipulation.*
Spatial Index	*To structure information in the environment using physical props and spatial layout.*
Material Cue	*To aid recognition and reduce search costs by exploiting graphical and material variables.*
Conversation Prop	*To provide shared representations that support discussion and make outcomes explicit.*
Social Currency	*To indicate work roles and relationships through ownership, possession and exchange.*

Table 6.1: Token Roles Derived from Coding of Interview Transcripts

To resolve these tensions, the category of "cognitive aid" was broken down according to an apparent three-way distinction between strategies for achieving cognitive support in the physical world. The first distinction made was based on handling – tokens can be used as "knowledge handles", reflecting their ability to stand for knowledge in the user's head, in a manner that permits the user to externally manipulate that knowledge in the world. This token role captures those aspects of "interaction device" related to the exploratory or experimental handling of tokens whilst thinking about their associated digital attributes. The second distinction made in the provisional "cognitive aid" category served to divide the remaining uses of tokens – those that did not directly involve handling – into categories reflecting exploitation of tokens' material and spatial affordances respectively. Tokens can therefore also provide cognitive support through their use as "material cues" or "spatial indices", drawing attention to both themselves and the things that they stand for through their meaningful physical appearance or spatial location respectively.

The original category of "social currency" also had a distinctive internal structure that led to its decomposition into two new categories. One of these retained the title of "social currency", but narrowed its definition to cover only those aspects of use related to token exchange. The newly created category of "conversation prop" captured those aspects of the previous definition of social currency related to social encounters around tokens that are not exchanged.

Finally, the remaining excerpts from the original category of "interaction device" were relabelled as "digital instruments" to specifically highlight the use of tokens as instruments for the inspection the manipulation of digital attributes. Table 6.1 presents a summary of the six resulting token roles. A more thorough discussion of each token role, accompanied by selected interview excerpts, is given in Appendix A.9.

Chapter 7

Elicitation of Experienced Benefits

The previous chapter described the deployment of a prototype TUI in a real office context, and presented qualitative judgements of the impact that the TUI had made on individual and collective work practices. In this chapter, I introduce a method for determining which characteristics of TUI functionality and usage contribute most towards the benefits derived from interactions with the TUI over time, and go on to present the results of the method's application to the users of my TUI prototype in its contextual deployment at HomeWatch.

7.1 The Need for Experienced-Benefits Elicitation

Despite the prevalence of assertions in the TUIs literature as to the benefits of tangible interaction, few publications report any kind of in-depth user testing or field deployment. Marshall (2007) draws attention to the fact that "empirical work and theoretical development have failed to keep up with the pace of technical development" in the area of TUIs, arguing that "theoretically-grounded accounts and empirically-based studies are now needed to understand better how tangible interfaces actually work". More generally, Rogers (2006) has argued that "more studies are needed of UbiComp technologies being used *in situ* or the wild – to help illuminate how people can construct, appropriate and use them".

To address concerns regarding the lack of data substantiating claims of tangible benefits, I devised a method for generating and evaluating benefit hypotheses, which attempts to bridge the divide between benefits motivated by theory on the one hand, and by user experiences on the other.

The method is based on the concept of *benefit factors*, that is, factors contributing to the overall benefits of using tangibles for peripheral interaction. This method has two sequential components: a *communication* stage, where benefit hypotheses are systematically generated both from established theory and from researcher reflection, and translated into terms comprehensible to users of the technology; and an *elicitation* stage, where these benefit factors are presented to users of the technology as props for the discussion of experiences and the rating of relative contributions. This is an application of the *graphic elicitation* method of Crilly et al. (2006), which articulates a case for using research diagrams as an interview stimulus, in order to guide discussion and derive feedback on researcher conceptualisations. The details of my *experienced-benefits elicitation* method are discussed next.

7.2 Communication of Experienced Benefits

A method was required that could be used to generate a range of potential benefits arising from the tangibility of the interaction objects used in peripheral interaction. A systematic approach should serve to highlight those areas that have not been addressed by theory, encouraging the researcher to formulate their own hypotheses concerning factors contributing to benefits in these unexplored areas of 'benefit space'.

Benefits can be experienced in a number of different ways, and expressed at multiple levels of abstraction. They can arise through the interplay of different aspects from one or many objects, just as multiple benefits can arise from any one part of an interacting system. A framework was required to focus attention on these competing factors, in a manner that encouraged systematic and reflective thinking.

Two distinct dimensions were chosen for the mapping of benefit space: the *domain of experience* in which benefit factors are realised, and the *mode of reflection* in which benefit factors are considered and expressed. This particular dimensionalisation was motivated by Husserl's phenomenology (Husserl, 1927), which recognises that reflection on experiences is the way in which we become conscious of phenomena:

> *Focusing our experiencing gaze on our own psychic life necessarily takes place as reflection, as a turning about of a gaze which had previously been directed elsewhere. Every experience can be subject to such reflection…*

It is this introspection by users of the system that will reveal their subjective judgements of benefit, with experience domains and reflection modes probing different areas of their experience in different ways. Each combination of experience domain and reflection mode is therefore called a *benefit area*. These are introduced briefly next, before going on to discuss how each area was used to generate benefit factors. These factors were then translated onto palm-sized cards giving names, descriptions and pictures, to be used in the elicitation exercises discussed at the end of this chapter. The designs of these cards are presented in Section 7.2.3.

7.2.1 Domains of Experience

The *interactional*, *cognitive*, and *social* domains of experience were adopted for the purpose of describing the nature of user benefits, loosely corresponding to various types of experience:

the mechanical and perceptual aspects of interacting with the system; the mnemonic, computational and interpretive aspects of interacting with the world; and the interpretive and communicational aspects of interacting with other people.

7.2.2 Modes of Reflection

Different modes of reflection were modelled as an abstraction hierarchy, in which each mode corresponds to an abstraction level that generalises over all that precede it. The lowest level is *ability*, corresponding to the human abilities that can be augmented or exploited by tangible interaction. Building on this is the *action* level, based on the interface-specific actions the user can perform by making use of their abilities. Next is the *appropriation* level, where users perform actions and assign meaning in a manner that is not part of the formal notation provided by the system. User actions and appropriations make use of abilities to support generic work at a higher level of abstraction, examined through the *activity* level. Finally, the *atmosphere* level accounts for global changes that occur over time as a result of the extended application of tangibles to support work activities.

7.2.3 Benefit Areas

The following sections describe the 15 different benefit areas, corresponding to the product of the three domains of experience with the five modes of reflection. Each area was used to prompt researcher consideration of the ways in which tangibles could provide benefits in that domain of experience, and expressed in that mode of reflection. For each area, the aim was to identify four particularly salient factors that could be seen to contribute to the overall benefits of using tangibles for peripheral interaction. This process generated 60 different such factors for the forthcoming comparison by users. This systematically generated, uniform framework was not expected to be complete, but it was intended to be comprehensive. By giving careful consideration to fifteen different benefit areas, all of the major potential benefits and ways of expressing them should have been covered, along with some minor benefit factors that have perhaps not been considered previously.

Each of the benefit factor names and descriptions given in the following sections are the translated benefits intended for assessment by participants in the elicitation session. The wording of each description begins with an indication of its parent reflection mode, with the terms "use" and "effect" substituted for the terms "appropriation" and "atmosphere" respectively, in order to improve clarity.

Interactional/Ability: Perception and Actuation

Two-Handed Dexterity
The ability to use all of the digits of both hands in a coordinated manner.

Bodily Awareness
The ability to sense the position of the limbs and their movements, with eyes closed.

Recognition
The ability to acknowledge that an object is familiar and has been perceived before.

Peripheral Vision
The ability to perceive objects lying outside the direct line of sight.

Cognitive/Ability: Attention and Spatiality

Attention Focus
The ability to use objects to focus attention on a task using those objects.

Attention Diversion
The ability to use objects to divert attention from a task not using those objects.

Spatial Memory
The ability to recall the locations of objects without looking.

Spatial Reasoning
The ability to mentally simulate changes in the arrangement of objects.

Social/Ability: Value and Persuasion

Information Value
The ability to assign value to objects based on the information they represent.

Social Value
The ability to assign value to objects based on potential use in social encounters.

Power
The ability to directly affect the way another group member carries out their work.

Influence
The ability to indirectly affect the way other group members carry out their work.

Interactional/Action: Interactivity and Control

Token Marshalling
The action of moving tokens between the interactive surface and the desktop.

Surface Glancing
The action of quickly and periodically glancing at the interactive surface.

Multi-Point Controlling
The action of using the fingers to move multiple tokens at a time.

Eyes-Free Controlling
The action of using the hands to interact, with visual attention elsewhere.

Cognitive/Action: Handling and Organisation

Playful Fiddling
The action of idly handling tokens in a non-purposeful manner.

Exploratory Handling
The action of manipulating tokens to uncover hard to visualise information.

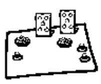

Desktop Structuring
The action of using physical props to structure information on the desktop.

Desktop Scanning
The action of visually scanning the desktop environment for particular tokens.

Social/Action: Conversation and Sharing

Conversation Focus
The action of using tokens as a means of guiding conversation.

Conversation Initiation
The action of using tokens as a means of starting conversation.

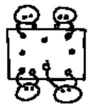

Token Pointing
The action of pointing to tokens as a means of conveying information.

Token Passing
The action of passing tokens to others as a means of conveying information.

Interactional/Appropriation: Convenience and Contrast

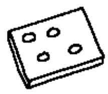

Interaction Focus
The use of token surface-presence as an indicator of interaction focus.

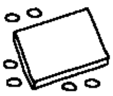

Interaction Readiness
The use of token-surface proximity as an indicator of interaction readiness.

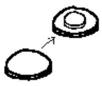

Physical Adornment
The use of token surfaces as places for the attachment of materials.

Graphical Annotation
The use of token surfaces as places for erasable annotation.

Cognitive/Appropriation: Materials and Space

Iconic Objects
The use of objects to signify meaning through similarity.

Symbolic Materials
The use of materials to signify meaning through arbitrary rules or conventions.

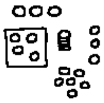

Relational Layout
The use of physical space to convey relationships between tokens.

Relative Location
The use of desktop areas and objects to give meaning to nearby tokens.

Social/Appropriation: Possession and Exchange

Rights
The use of token possession as a social indication of rights.

Responsibilities
The use of token possession as a social indication of responsibilities.

Agency
The use of token possession as an indication of the right to act of behalf of the owner.

Expectations
The use of token possession as an indication of the expectations held by the owner.

Interactional/Activity: Input and Output

Digitising
The activity of taking a concept in the head and recreating it on a computer.

Updating
The activity of making incremental changes to information stored on a computer.

Exploring
The activity of rapidly trying out many different information states.

Monitoring
The activity of periodically checking a computer for other people's updates.

Cognitive/Activity: Offloading and Cueing

Externalising
The activity of creating external cues to guard against memory failure.

Cognitive Offloading
The activity of reducing memory load by manipulating artefacts in the world.

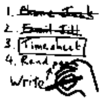

Tracking
The activity of reducing memory load by tracing changes in the world.

Reminding
The activity of referring to external cues as a means of prompting the memory.

Social/Activity: Ideation and Evaluation

Brainstorming
The activity of generating a wide variety of ideas for later analysis.

Planning
The activity of formulating things to be done into a specific course of action.

Decision Making
The activity of evaluating ideas and deciding upon appropriate actions.

Reviewing
The activity of looking back on the work that has been achieved.

Interactional/Atmosphere: Productivity and Affect

Efficiency
The effect of achieving the same amount of work in less time.

Ergonomics
The effect of having tools that better fit user capabilities and the structure of tasks.

Engagement
The effect of having tools that sustain users' attention and encourage interaction.

Enjoyment
The effect of having tools that are fun to use and fun to share with others.

Cognitive/Atmosphere: Quality and Availability

Accuracy
The effect of increasing the degree to which information reflects the work situation.

Timeliness
The effect of increasing the degree to which information is available on time.

Quantity
The effect of increasing the quantity of information that exists.

Accessibility
The effect of increasing the degree to which existing information can be retrieved.

Social/Atmosphere: Teamwork and Integration

Collaboration
The effect of increasing the amount of jointly completed work.

Coordination
The effect of increasing the amount of work done to suits others' work schedules.

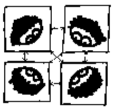

Awareness
The effect of increasing the degree of mutual awareness about the work of the group.

Cohesion
The effect of increasing the level of social contact between group members.

7.3 Elicitation of Experienced Benefits

This section describes the benefits elicitation method as it was applied in the HomeWatch field deployment presented in Chapter 6. The procedure is presented first of all, followed by the motivations behind the technique. Next, an analysis of user responses is presented, and used to construct a characterisation of the design space based on their experiences.

7.3.1 Procedure

The benefits elicitation process consists of two stages: ranking and rating. The rationale is to give users two opportunities to consider each benefit factor and to compare it to others, as well as giving the researcher the opportunity to analyse the correspondence between these two datasets.

Ranking Exercise

In the ranking exercise, comparisons between benefit factors are made in the sets of four corresponding to each benefit area, providing users with an in-depth opportunity to clarify what is meant by each benefit factor and to make a small number of comparisons in terms of the relative impact that the four factors had made on their experiences of TUI use (in this case, the prototype system presented in Chapter 5). The initial setup, with the benefit factor cards arranged according to benefit area in 15 piles of four, is shown in photograph (a) of Figure 7.1.

The ranking of the four benefit factor cards from each set is achieved by arranging each set of four cards vertically on the tabletop, with factors of equivalent impact placed side by side within this vertical arrangement. While participants consider each set in turn, they are asked to think aloud about the experiences they are drawing on in making their assessment. Each spatial arrangement can be recorded conveniently by photograph before moving on to the next. Some example spatial rankings from probe participants' elicitation exercises (and their equivalent numeric rank) are shown in photograph (b) of Figure 7.1.

Figure 7.1: Photographs of Benefit Elicitation Exercise. (a) Initial preparation of benefit factors into related sets of four; (b) spatial rankings of benefit factors in three benefit areas, with corresponding numeric ranks; (c) use of paper plates as a tangible rating scale; (d) final ratings of benefit factors.

Mode of Reflection	**Question Asked to Elicit Feedback on Associated Benefit Factors**
Ability	To what extent does this kind of interface *exploit* these abilities?
Action	To what extent does this kind of interface *enable* these actions?
Appropriation	To what extent does this kind of interface *encourage* these uses?
Activity	To what extent does this kind of interface *support* these activities?
Atmosphere	To what extent does this kind of interface *achieve* these effects?

Table 7.1: Question-based Elicitation of Benefit Factors

Rating Exercise

Whilst for the ranking procedure participants are asked to reflect on their experiences of using a particular prototype, in the rating exercise, participants are asked to project their experiences onto how they see each factor contributing to the overall benefits of using tangibles for peripheral interaction, assuming that a similar interface were to be deployed in a similar context, but with complete and permanent adoption in terms of personnel and timescale. This adjustment is made to capture the perceived benefits of interaction experiences that would only arise out of ongoing and widespread use, such as those based on network externalities. The means of rating each benefit factor is a seven-point interval scale; in terms of the contribution of each benefit factor to the overall benefits of peripheral interaction with tangibles, this scale ranges from "contributes very little", to "contributes very much". The 12 cards from each "mode of reflection" are shuffled and presented sequentially in these sets of 12, until the participant has rated all 60 benefit factors. The question asked of the interviewee varied between these modes of reflection, as shown in Table 7.1.

Each point on the scale is represented by a paper plate labelled with a number from one to seven – these are lined up from left to right in front of the participant. As they consider each card they turn, they are asked to place it on one of the plates such that the names of any cards already on the plate remained visible. Placing a card on the plate labelled X is equivalent to assigning the corresponding benefit factor a rating of X on the 7-point interval scale. Participants have the freedom to move cards between plates at any time. Photograph (c) of Figure 7.1 shows this rating procedure in action.

After all of the cards of a perspective have been assigned to plates representing their rating, they are lined up vertically above the corresponding plate. At this point, the participant is given the freedom to change the rating of previously ordered cards by moving them from one line to another, as well as changing the rating of just seen cards by moving them from their plate onto a different line. The procedure is repeated for all five modes of reflection, at which point participants are given a final chance to rearrange the cards into their overall assessment of benefit contribution. Photograph (d) of Figure 7.1 shows the result of such a rearrangement.

7.3.2 Motivations

The primary motivation for using a card-based elicitation, rather than a questionnaire-based structured interview, is to reduce the effects of the cognitive anchoring and adjustment bias (Buehler et al., 2002). The manifestation of this bias in questionnaires occurs when the existence of an anchor, in the form of the rating given to a questionnaire item, results in subsequent ratings being adjusted insufficiently relative to this anchor, meaning that ratings are affected by the order in which items appear. This is compounded by the permanence of questionnaire-based ratings, the difficulty in directly comparing items that are distant from one another in the presentation order, and the difficulty in seeing commonly rated items as a group. The card-based elicitation method addresses all of these issues.

First of all, the two-pass approach means that participants will be able to make better initial ratings in the second stage, since they will have already seen the whole set of factor cards in phase one. Whilst anchoring and adjustment is the natural cognitive strategy by which relative comparisons are made, this approach reduces that likelihood that an inappropriate initial anchor will be chosen. It also provides two distinct data sets for an analysis of correspondence between experienced benefits and projected benefits, comparing the initial ranks of each set of four cards with the ranks of their subsequent rating (see Section 7.3.3).

Secondly, the presentation of commonly rated items as a group increases the conviction with which decisions can be made, since each comparison is made against not one but multiple points of reference. Either the new card 'fits' the existing group, or suggests comparison with the group either above or below in the rating scale. It also allows the direct comparison of items from multiple domains of benefit and modes of reflection that may otherwise have been presented far apart in a sequential questionnaire.

Finally, the reconfigurability of cards' spatial positions, and hence ratings, means that decisions are easily reversible, and that difficult-to-rate cards can be deferred until more decisions have been made, increasing the number of available points of reference against which to compare. In particular, it allows users to see the gestalt of their decisions and 'smooth' the distribution of their responses once they see the big picture. At this stage, any changes are made in the context of all previous ratings and re-ratings, giving the best possible opportunity to obtain a reliable indication of participants' true assessments[1].

7.3.3 Analysis of Initial Ranking Exercise

One of the advantages of the 'ranking then rating' procedure is that the two resulting datasets can be compared. However, the unknown nature of the rating distribution, and

[1] In practice, probe participants' decision-making behaviour in the elicitation sessions was consistent with that expected. Their performance, in terms of bimanual and epistemic actions within a gestalt view of previous yet provisional decisions, suggests that translating a structured interview into card format offers a number of benefits in terms of review and reflection, at the expense of time spent on interview preparation and on the interview procedure. This trade-off is analogous to the trade-off between the strength of TUIs and the generality of GUIs – in a sense, the card-based elicitation of experienced benefits can be seen as exploiting many of the same benefits of tangibility as TUIs.

the ordinal and interval nature of the data, called for a non-parametric statistical test of the association between the two sets of responses. The ratings were first translated into ranks, with adjacent tied ranks in either data set assigned the mean value of the ranks that were tied – e.g. ratings of [5, 6, 6, 7] would become ranks of [1, 2.5, 2.5, 4]. Due to the high incidence of ties in both data sets, a measure of association was required that was weakly monotonic, i.e. as x increases, y either increases or stays the same, and vice versa. This would not penalise participants – in terms of reducing the measured degree of association – who either differentiated in the rating exercise between factors that they tied in the ranking exercise, or who gave the same rating to factors of differing initial rank. It would, however, highlight inconsistencies between the data sets. Goodman and Kruskal's *gamma* coefficient was chosen for this purpose (see Appendix A.8 for a description).

Given that the two exercises were completed sequentially and were asking similar questions, it was expected that the majority of comparisons would yield positive *gamma* coefficients (showing that the ordering of benefit factors in the ranking exercise was consistent with their relative ordering in the rating exercise) . Of the 15 four-card sets, the number of sets in which all cards were ranked or rated as equivalent by P1, P2 and P3 were 2, 1 and 3 respectively, resulting in undefined *gamma* coefficients (due to division by zero). Given that the aim of the comparison was to evaluate consistency, a new function was defined that converted *gamma* coefficients to a measure of domain-specific consistency.

$$consistency = \begin{cases} gamma & \text{if } gamma \text{ is defined} \\ 1 & \text{if } gamma \text{ is undefined because all ranks are rated equal} \\ 0 & \text{otherwise} \end{cases}$$

This *consistency* function maps undefined *gamma* coefficients to 1 if different ranks mapped to the same rating, i.e. the ranking could be viewed as occurring within a rating interval, and 0 otherwise. An inconsistent ranking/rating combination within a benefit area was defined as any *consistency* value of 0 or less. The results are shown in Table 7.2.

The low proportion of inconsistent responses by all users is in accordance with the view that users' experience with the system should have a strong influence on their projection of benefits onto any similar system and context. However, it also reflects on how users perceived the relationships between the cards in each set. For instance, the social/action factors of *token passing*, *token pointing*, *conversation initiation*, and *conversation focus* were difficult for users to distinguish between, in terms of their individual contributions to overall benefit, as were the cognitive/activity factors of *externalising*, *tracking*, *cognitive offloading* and *reminding*. In a number of situations, requests for clarification at the rating stage led to responses that were revised from the earlier ranking. This is a positive consequence, because it means that the final ratings were more accurate than they would have otherwise been. It occurred most frequently with P2, who 'promoted' *recognition* and *information value*, and 'demoted' *power* and *attention diversion*, based on interviewer clarification.

Domain of Experience	**Mode of Reflection**	*consistency* (******gamma* **undefined)**		
		P1	P2	P3
Interactional	Ability	1.00	-0.20	1.00
Interactional	Action	0.60	1.00	1.00
Interactional	Appropriation	1.00	1.00	0.00
Interactional	Activity	1.00	0.00	0.50
Interactional	Atmosphere	1.00	1.00	1.00
Cognitive	Ability	0.20	-0.20	0.67
Cognitive	Action	1.00	0.50	0.50
Cognitive	Appropriation	-1.00	0.33	0.60
Cognitive	Activity	1.00*	1.00	0.00*
Cognitive	Atmosphere	1.00	0.50	0.50
Social	Ability	0.20	-1.00	1.00
Social	Action	1.00*	0.00*	0.00*
Social	Appropriation	1.00	1.00	1.00*
Social	Activity	1.00	1.00	1.00
Social	Atmosphere	1.00	1.00	1.00
Proportion inconsistent (*consistency* ≤ 0)		0.07	0.33	0.27

Table 7.2: HomeWatch Results: Association between Rankings and Ratings

7.3.4 Analysis of Final Rating Exercise

Inspection of the distribution of ratings assigned by each participant revealed that they were neither normally-distributed, nor of the same distribution. Since this meant that individual comparisons of ratings would be susceptible to bias introduced by this variation in rating distributions, the raw ratings were therefore converted to ranks for comparison. However, this resulted in a large number of ties due to the relatively small number of rating levels compared with the number of rated benefit factors.

To resolve as many of these ties as possible, an assumption was made that the initial rank of a benefit factor card would provide a reasonable approximation to its relative position within its rating level. Hence within each subset of factors rated equivalently, those factors placed first in their set of four would be ranked higher than all those placed second – in the same set of four, or any other. A final rank ordering for each participant was thus produced, as shown in Figure 7.2.

In order to examine the uniformity or otherwise of the benefit space mapped out by the dimensions of experience domain and reflection mode, the aggregate rank of the four factors in each benefit area was calculated for each benefit area and for each participant. A visualisation was created which could highlight any differences both between benefit areas of participants and between participants themselves, as shown in Figure 7.3. This visualisation immediately draws attention towards a number of patterns. P1 and P2 generally had similar overall ratings of benefit areas, except in two cases. The general agreement between these two can be seen as a consequence of their similar technical orientation and their shared

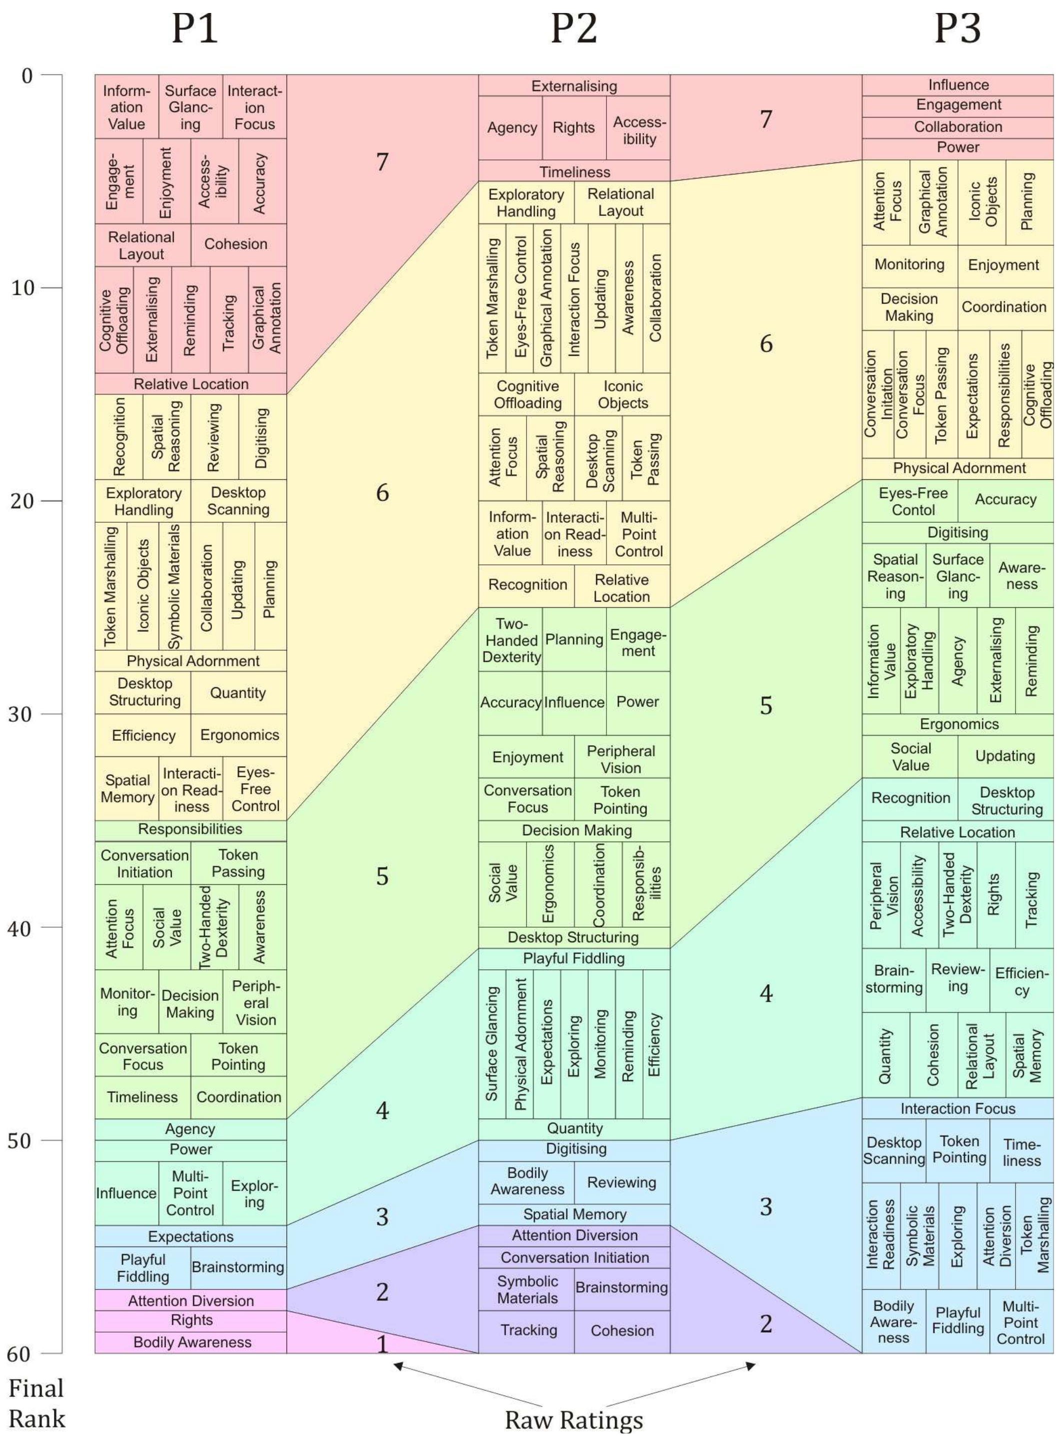

Figure 7.2: HomeWatch Results: Overall Rankings of Benefit Factors

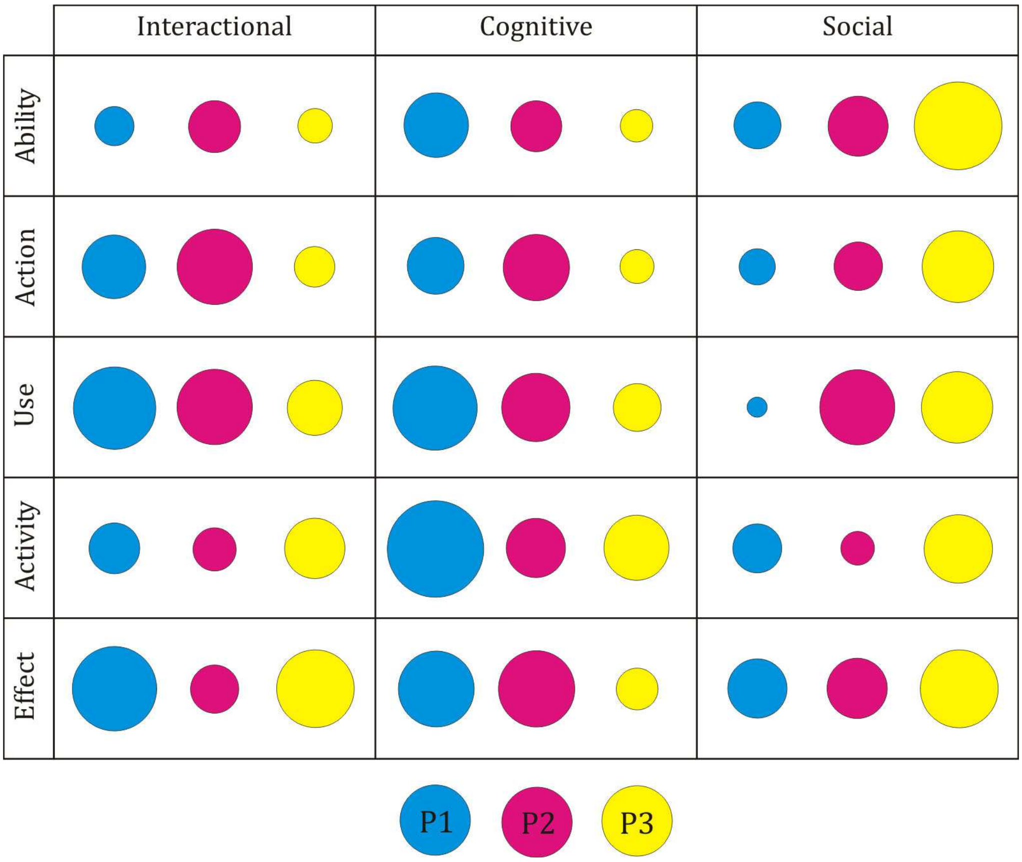

Figure 7.3: HomeWatch Results: Comparison of Benefit Areas. Bigger circles correspond to greater benefit derived from that benefit area (represented by table cell), according to that user (represented by circle colour). Cells allow comparison within benefit areas across participants, rows allow comparison within modes of reflection across domains of experience, and columns allow comparison within domains of experience across modes of reflection.

experiences. For seven out of ten of these areas in the interactional and cognitive domains of experience, though, P3's aggregate ranks were lower than both P1 and P2. However, for four out of five of the social benefit areas, P3's aggregate ranks were the highest. Her preference for social areas of benefit, rather than interactional or social, are again likely to be a function of her work role as the office manager. This requires a high level of organisation and significant time away from her desk, interacting with other office members, in meetings, or on errands. P3 also displayed reasonably uniform aggregate rankings for social areas of benefit, as did P1 and P2 for cognitive benefit areas. No participant displayed any such strong, consistent preferences for any one mode of reflection. This too makes sense – the mode of reflection should not significantly affect the perceived benefit of the experience, or imagined experience. Finally, P3's aggregate ratings for interactional benefit areas showed a steadily increasing trend moving from abilities to effects. This can be explained by the fact that her role was non-technical, and as such she would be more likely to perceive benefit in the effects of interaction, rather than the mechanisms by which they were achieved.

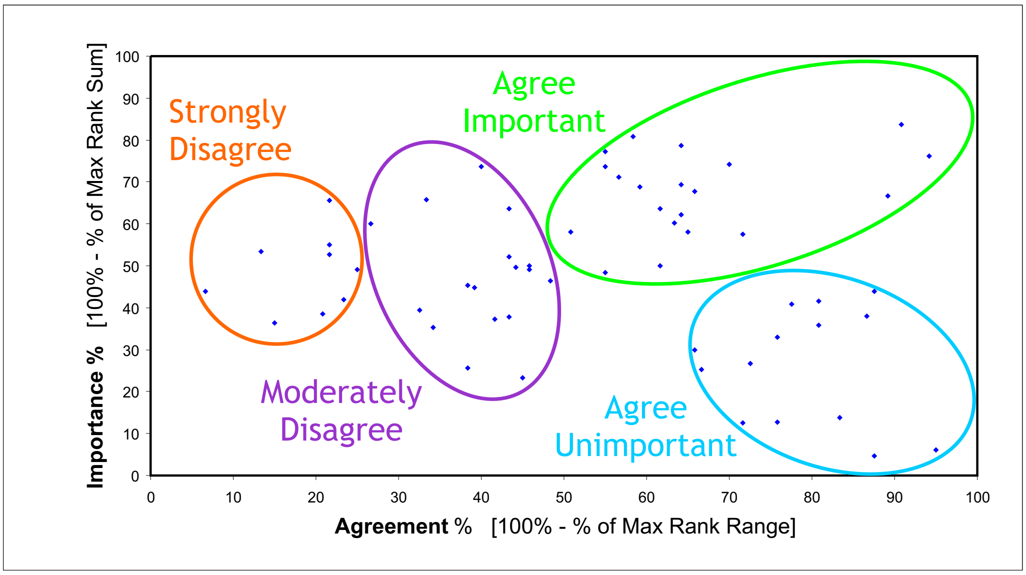

Figure 7.4: HomeWatch Results: Scatter Plot of Aggregate Rank vs. Rank Range

Overall, the visualisation of benefit areas in Figure 7.3 shows that the mapping of benefit space was a reasonably balanced sampling. Whilst it highlighted a number of response patterns, there were no areas that performed significantly better or worse than any other, taking into account all three sets of responses. The visualisation also does not provide a clear depiction of how the aggregate ranks of each benefit factor correspond to its range of ranks, i.e. the level of agreement, which could be extracted from the final ranking diagram of Figure 7.2. To address this issue, a scatter plot of *importance* (derived from the sum of each participant's rank for a particular factor), versus *agreement* (derived from the range of ranks across participants for a particular factor) was created, as shown in Figure 7.4. Four areas of the scatter plot have been highlighted, corresponding to the following categories of benefit contribution: *strongly disagree*, *moderately disagree*, *agree unimportant*, and *agree important*. Aside from the unremarkable benefit factors falling under the category of "moderately disagree", these categories will now be used to structure the analysis of participants' responses.

Strongly Disagree

This category contains the nine factors with the largest differences in assigned rank. The most controversial use of tokens was as an indicator of *rights*, receiving the maximum rating from P2, the middle rating from P3, and the lowest rating from P1. The root of the disagreement lay in participants' interpretation of how rights would be realised, with P1 viewing them as a tool of constraint, in contrast with P2's view as a tool of enablement. The act of transferring rights from one person to another by token passing creates a relationship of *agency*, whereby the receiver of the token is understood to be able to act on behalf of its owner. The relationship between the two was observed by all participants, who similarly disagreed about its potential benefit.

Related to the expressions of rights and agency are the abilities to exert *power* and *influence*. Neither P1 nor P2 rated these highly, yet they both received maximum scores from P3. This distinction is likely a function of the types of work engaged in by the participants – whereas P1 and P2 were a tightly integrated team of two who essentially managed their own work, P3's role as office manager frequently involved the coordination of all office members. The use of physical tokens as a means of getting people to do things that they need to do, but perhaps might otherwise forget or not want to do, would be a welcome prospect for someone in such a role.

The roles of participants also had a bearing on the granularity and number of their tasks. Whereas P1 and P2 had between ten and twenty independent, medium-term things to do at any one time, P3 had a small number of areas of responsibility that each contained many very small work items such as emails and phone calls. Hence for P1 and P2 the interactive surface provided an *interaction focus* that was of the highest benefit, since it conveyed all of the things that they should be working on. However, the number and transient nature of P3s work items meant that the interactive surface did not provide a single point of focus for task management. It also meant that *token marshalling* was not as important for P3, since she didn't have as many tokens to manage. The *timeliness* with which information was made available to others was similarly not seen as a particular benefit by P3, perhaps due to the fact that her tokens were too high-level to prompt action on any particular work item associated with them, but not physically represented by them.

P1 gave the highest rating to *tracking* and *cohesion*, which both received low and middle ratings from P2 and P3 respectively. The relatively high rating of tracking by P1 may be due to the fact she saw it as part of a group along with the other aspects of external cognition, and didn't consider that whilst tokens provide ample focus for externalising, reminding, and cognitive offloading, their use as a medium for tracing changes over time is relatively limited. Alternatively, it could derive from different interpretations of the nature of tracking; P1, for example, made extensive use of the task timer, which could be seen as a way of tracking the time spent on tasks. Finally, the benefit arising from increased cohesion – using tokens to encourage social interaction – depends on a number of factors including personal preference, social structures and token usage patterns.

Agree unimportant

This category contains the fourteen factors with relatively small differences in assigned rank, combined with relatively high aggregate rank (less important). There are a number of ways of explaining why these factors were given low ratings by all participants. For some factors – *bodily awareness* and *two-handed dexterity* – my communication of the concepts, in hindsight, may have been unclear. The bodily awareness card used the phrase "with the eyes closed" rather than the more appropriate "without visual attention", perhaps explaining why it received much lower ratings than the associated factor of eyes-free controlling, which relies upon it. Similarly, the concept of two-handed dexterity conflated the related but distinct concepts of bimanual skill and manual dexterity; the fact that the prototype interface requires a high level of bimanual skill but a low level of manual dexterity could therefore have led to participant uneasiness in rating this benefit factor.

Token Role Category	Benefit Factors Central to that Category
Digital Instrument	updating, eyes-free controlling
Knowledge Handle	information value, exploratory handling, cognitive offloading
Spatial Index	externalising, relative location
Material Cue	graphical annotation, physical adornment, iconic objects, recognition
Conversation Prop	planning, spatial reasoning
Social Currency	token passing, responsibilities

Table 7.3: HomeWatch Results: Distribution of ”Agree Important” Benefits

A number of factors were likely to provoke assessment difficulties – *quantity* and *ergonomics;* others were perhaps inappropriate for the kind of interaction and media involved – *exploring* and *brainstorming;* whilst some were considered unnecessary – *attention diversion* and *playful fiddling*. For the remainder, however, it could just be that the contribution was relatively unspectacular: *spatial memory* is used to a limited extent for the location of both the Powermate and for tokens; *peripheral vision* is used in an indirect way, but is better expressed through concepts such as recognition and reminding; and *desktop structuring*, *social value*, and *token pointing* are similarly all related to the more strongly perceived benefits of relational layout, token passing, and conversation focus respectively.

Agree important

This category contains the 20 factors with low aggregate rank (more important), and low rank range, representing those factors that participants agreed were most important. There were five atmospheric effects in this category, which I hypothesised might be related in a feedback loop. The highest total belonged to *engagement* – "tools that sustain users' attention and encourage interaction". This makes them "fun to use and to share with others" – my definition of *enjoyment*, which encourages *collaboration* by bringing people together, which results in better mutual *awareness* when they are apart. This heightened awareness increases *accuracy* when making judgements about team-related business, which motivates a higher level of *engagement* when using the tool in future, thus repeating and reinforcing this 'cycle of engagement'. The querying of this hypothesis is presented in Section 7.5.

The remaining 15 non-effect benefit factors are spread amongst the six token roles presented in Section 6.4, as shown in Table 7.3. Token roles can be seen as radial categories with central members, which motivate the outer members through chains of association[2]. These central members should be chosen based on benefits that are clearly distinguishable from other benefits, intuitively understandable based on experience, and of general appeal to as high a proportion of users as possible. In the context of this fieldwork, the non-effect, "agree important" factors in Table 7.3 meet these criteria, and were therefore used as central members of categories that were subsequently expanded to incorporate all non-effect factors. Effects

[2]Based on Lakoff's (1987) concept of radial categories from the book *Women, Fire, and Dangerous Things.*

were not included in this categorisation process since they arise from the interplay between multiple factors drawn from multiple roles.

Figure 7.5 presents the result of this category expansion. The central members of each token role are contained within a rounded box, with motivating links drawn between both them and peripheral members of the category. The font size and font darkness of each factor are a function of its aggregate rank and rank range respectively[3].

The accommodation of all non-effect factors in this breakdown suggests that the six token roles identified during interview-transcript analysis provide a comprehensive structuring of token uses. In this diagram of token use, each factor belongs to a single category, and no links between categories are shown. However, the purpose of constructing the model was not to create an accurate representation of reality, but to help identify themes of benefit that cut across all categories of use. Some factors contributed to the overall benefits of tangibles more than others, in the experiences of the three system users, and it is this juxtaposition of central and peripheral category members that advances the analysis from factors contributing towards the overall benefits of peripheral interaction, to an identification of the essential nature of the benefits themselves. This concern with the essence of experience is rooted in phenomenology, and its application to the benefits experienced in the technology probe deployment makes it the natural expression of the design outcome of the elicitation process – a phenomenologically-informed account of the essential qualities of peripheral interaction.

7.4 Essential Qualities of Peripheral Interaction

Figure 7.5 serves to emphasise the differences between the six fundamentally different roles that tokens can play in peripheral-interaction design. However, it also provides an opportunity to look for commonalities – themes of benefit that cut across role boundaries and define the *essential qualities* of peripheral interaction that distinguish it from other interaction paradigms. The distinction between the roles of tokens and the essential qualities of peripheral interaction with tokens is as follows:

- **Token roles in peripheral interaction** describe the *various forms of engagement* users can have with tokens, providing a vocabulary for use in both the formative design and summative evaluation of token-based interfaces.
- **Essential qualities of peripheral interaction** describe the various characteristics of tokens that encourage an *episodic kind of engagement*, providing both design rationale and design prescription based on interaction context.

The following sections describe the results of this analysis, highlighting the essential qualities that are found in all of the six token role categories, represented in each by relatively important benefit factors as agreed by the probe participants.

[3]Font sizes in the range [10,20], from smallest to greatest aggregate rank; HSV brightness levels [0, 60], from smallest to the greatest rank range.

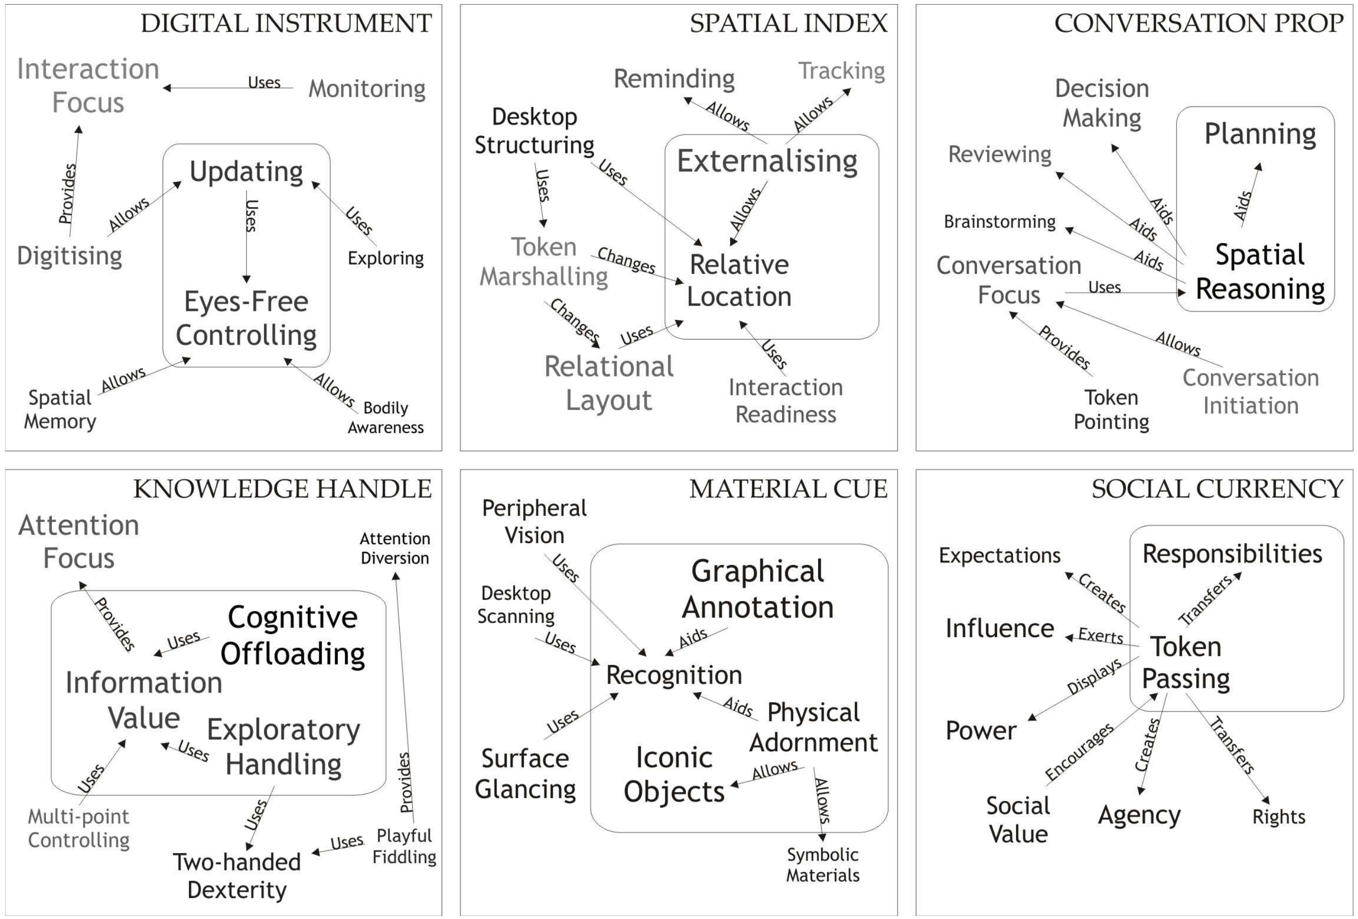

Figure 7.5: Analysis of Results: Structures of Token Role Categories. Each box represents a token role, which contains the benefit factors most closely associated with that category of use. Larger, darker text indicates greater importance and agreement with respect to the results of the elicitation exercise at HomeWatch. "Central" members of token-role categories are contained within rounded rectangles.

7.4.1 Frame of Reference

This common theme is about the local environment being rendered meaningful – whether the interactive surface, desktop or meeting table – in order to lower the barriers to interaction whenever it needs to occur. However, this theme is not about interaction, it is about *anticipation*. If a user believes that externalising information will be useful, they will be more likely to 'tokenise' their environment in anticipation of its use. Tokens used in such a way – spread around the local environment – provide benefits even when they are not being considered directly, by virtue of their future interaction potential. This observation leads to the first essential quality of tangible peripheral interaction: tangibles provide a FRAME OF REFERENCE for potential digital interactions via the interactive surface, physical interactions with desktop information structures, and social interactions with other people. They imbue the user's working habitat with meaning that exists, for the most part, on the periphery of their workspace and attention, providing a salient external structure that reflects their internal understanding of the past, present and future work situation.

Each token role had one or more benefit factors relating to tokens that are not the user's primary, or even secondary concern – just simply visible and accessible in the local environment:

1. DIGITAL INSTRUMENT: Providing an *interaction focus* – the benefit of simply having tokens present on the surface, even if they aren't currently being used.
2. KNOWLEDGE HANDLE: Performing *cognitive offloading* for the purposes of manipulating that information either immediately, or at some future point in time.
3. SPATIAL INDEX: Making use of the richly meaningful space of the desktop environment as a medium for *externalising* information.
4. MATERIAL CUE: The augmentation of tokens by *graphical annotation* and *physical adornment*, as a persistent indicator of token identity.
5. CONVERSATION PROP: The arrangement of tokens on a meeting table as a provision for *spatial reasoning* in the context of a discussion.
6. SOCIAL CURRENCY: The use of tokens to define, reflect and reinforce social relationships, such as *power* hierarchies, in between social encounters.

7.4.2 Focus of Attention

This common theme is about the use of tokens to draw attention towards the various kinds of information they represent, whether through nudging in a digital interaction, through handling and rearrangement in epistemic manipulation, or through deictic references in conversation. The ability of tokens to act as a passive FRAME OF REFERENCE, ready to support physical and cognitive engagement, leads to the second essential quality of peripheral interaction: tokens provide a FOCUS OF ATTENTION within the locus of users' extended visual and manual reach that acts as a signifier for the digital, mental, and social information they represent, as well as providing indices into users' episodic senses of memory relating to previous acts of engagement with those tokens.

This quality complements the passive structuring of information in the environment, which is of little benefit if it is never used. Just as each of the six token roles exhibited what we might call framing behaviour, so too do they contain behaviour that utilises the created frames in contextualising and focusing attention on particular points of interest:

1. DIGITAL INSTRUMENT: Using the interaction focus to concentrate on the token whose digital content requires *updating*.
2. KNOWLEDGE HANDLE: Using cognitive offloading with tokens to focus attention on the *information value* they represent.
3. SPATIAL INDEX: Using the externalisation of information to draw attention towards the meaning of tokens' *relative location* and *relational layout*.
4. MATERIAL CUE: Using graphical annotation and physical adornment to make tokens stand out as salient features of the environment, aiding *recognition*.
5. CONVERSATION PROP: Using the arrangement of tokens for spatial reasoning as a means of drawing attention through *token pointing*.
6. SOCIAL CURRENCY: Using physical representations of power and other relationships to focus attention on social contracts such as *responsibilities* and *agency.*

7.4.3 Flow of Activity

This common theme is about the use of tokens to ease the flow of thought, action, and information in the context of interactions with the system, the desktop environment, and other people – with users attending to tokens only peripherally as they perform their work activities. This concept of 'flow' is related to Csikszentmihalyi's (1975) use of the term, in which "action-awareness merging" is experienced based on a "sense of personal control" and "direct and immediate feedback". It leads to the third essential quality of peripheral interaction: tangibles support the FLOW OF ACTIVITY in which users are engaged, with the various aspects of external cognition supporting the flow of thought, the perceptual and kinematic advantages of using physical tokens supporting the flow of action, and the use of these tokens as persistent, portable, and personal handles to information supporting the transfer of information between users and their systems, between users and their local environment, and between one another.

The two qualities presented thus far represent opposite ends of the attention spectrum. However, most work activities are conducted in between the two extremes, with different tokens passing in and out of focus as users go about their work. This is also reflected in the structure of the six token roles:

1. DIGITAL INSTRUMENT: Performing updating by *eyes-free controlling*, to ease the flow of information from the user to the system.
2. KNOWLEDGE HANDLE: Using the information value of tokens during *exploratory handling*, using tokens for the fluid manipulation of ideas.
3. SPATIAL INDEX: Using relative location and relational layout to ease information structuring by *token marshalling*, and to support *reminding*.
4. MATERIAL CUE: Using the superior recognition of augmented tokens to support fast, frequent *surface glancing* that refreshes mental situational models.
5. CONVERSATION PROP: Using token pointing, and shared representations in general, for the exploratory allocation and discussion involved in *planning*.
6. SOCIAL CURRENCY: Using tokens to ease the communication of social contracts, such as responsibilities and agency, through *token passing*.

7.4.4 Freedom of Appropriation

A final essential quality of peripheral interaction, one which underlies all the others, is FREEDOM OF APPROPRIATION. The natural affordances of physical objects, given sufficient freedom from spatial syntax, encourage users to make their own sense of the world using tokens and the tools that come with them: markers, materials, containers, etc. Without such freedom, users would lose the ability to create external notations of cognitive or social structures. Whilst regular graphical user interface applications have been created to address cognitive and social needs, no graphical interface has the ability to provide such multifaceted information structures in as convenient a medium – portable, paintable, graspable, throwable, and so on. These four qualities – FRAME OF REFERENCE, FOCUS OF ATTENTION, FLOW OF ACTIVITY, and FREEDOM OF APPROPRIATION – define the fundamental benefits of peripheral interaction with tangibles.

The relationship between token roles and essential qualities is shown in Figure 7.6. Note how the FREEDOM OF APPROPRIATION quality underlies both the remaining three essential qualities, as well as providing the basis for the various token roles displayed on the left-hand side. Without this freedom, users would be constrained to interact with tokens in the limited manner determined by the formal syntax of the system. With this freedom, users can selectively and fluidly engage with tokens in an episodic manner, drifting between their framing and focusing qualities on the one hand, and their various usage roles on the other, all according to the FLOW OF ACTIVITY.

7.5 Validity of Results

To distil the main analytic concerns of peripheral interaction in a concise summary of fieldwork results, I created a synoptic diagram that depicts the relationships between the benefi-

FREEDOM OF APPROPRIATION	FRAME OF REFERENCE	FOCUS OF ATTENTION	FLOW OF ACTIVITY
Digital Instrument	interaction focus	updating	eyes-free controlling
Knowledge Handle	cognitive offloading	information value	exploratory handling
Spatial Index	externalising	layout & location	marshalling & reminding
Material Cue	annotation & adornment	recognition	surface glancing
Conversation Prop	spatial reasoning	token pointing	planning
Social Currency	power	agency & responsibilities	token passing

Figure 7.6: Analysis of Results: Relationship between Roles and Qualities. Three of the "essential qualities" (blue) cut across all "token roles" (yellow), which in turn are the result of the fourth essential quality: freedom of appropriation.

cial effects of interaction, the essential interaction qualities that give rise to those effects, and the roles of tokens that encourage interaction. The result is shown in Figure 7.7.

To derive this diagrammatic structure, I analysed the six token roles for their defining characteristics. I identified two independent dimensions that give structure to an "abstract usage space" for tangibles: a *spatial–material* dimension that distinguishes between fundamental aspects of physicality, and an *interaction–interpretation* dimension that distinguishes between fundamental types of activity. In the synoptic diagram, these dimensions are represented by two orthogonal axes that divide the plane into four quadrants: spatial-interaction, spatial-interpretation, material-interaction, and material-interpretation. I then made a secondary distinction between the individual and social 'spheres' of experience, represented in the diagram as concentric elliptical zones with individual activities embedded within the social. Each combination of diagrammatic quadrant and zone uniquely defines at most one token role, represented by green ellipses in Figure 7.7. Two descriptive tags, denoting representative uses of tokens in each role, were then placed next to the corresponding green ellipse as additional elements of the concise token role definitions. These can be read according to their diagrammatic context as follows:

1. a *digital instrument* is about individual spatial interaction, used to glance and update;
2. a *knowledge handle* is about individual material interaction, used to offload and explore;
3. a *spatial index* is about individual spatial interpretation, used to organize and remind;
4. a *material cue* is about individual material interpretation, used to customise and recognise;
5. a *conversation prop* is about social spatial interaction, used to reason and refer;
6. a *social currency* is about social material interpretation, used to value and exchange.

Next, I arranged the four essential qualities of peripheral interaction in the centre of the diagram as red rectangles, such that each quality was adjacent to the token roles that predominantly contributed towards it. These associations are made explicit by the use of dotted-line boundaries to represent dependency by containment. Arrows between the red rectangles indicate relationships of support between interaction qualities. Finally, the five most beneficial "effects" of system use according to the probe participants were added to the diagram as orange rounded rectangles, and linked together according to the 'cycle of engagement' that I hypothesised in Section 7.3.4. Each effect was located in the diagram next to the token roles that I believed contributed most towards it. The centre of the diagram – enclosed by the four essential qualities of peripheral interaction and at the intersection and midpoint of each of the underlying axes – is the effect of *engagement*. This is fitting, since engagement was the most highly rated effect of interaction with my prototype TUI by probe participants, and it can be seen as equally relevant to the spatial and material concerns of both interaction and interpretation. However, it is also surprising that a system designed to support low-attention, peripheral interaction should be prized as a tool that "sustains users' attention

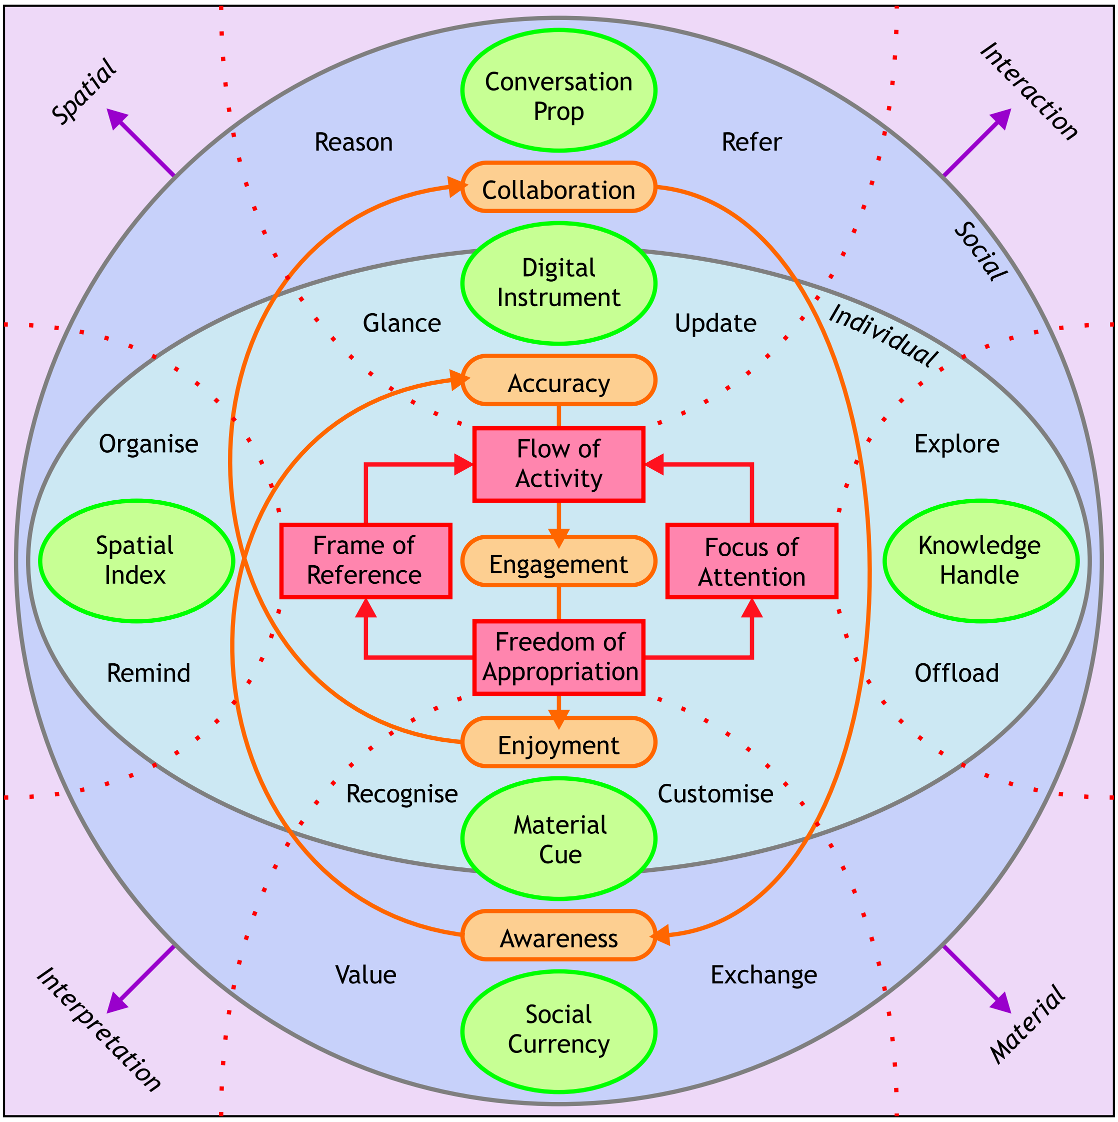

Figure 7.7: Validity of Analysis: Synoptic Summary for Member Checking. Two dimensions, interaction–interpretation and spatial–material, define four quadrants used to partially define token roles (green ellipses), further defined by their location within the individual or social 'spheres' of activity. The four interaction qualities (red rectangles) are arranged around the centre of the diagram next to their closest related token roles, surrounding the most highly rated beneficial effect of peripheral interaction: engagement. Other beneficial effects (orange rounded rectangles) connect to this in a positive feedback cycle.

and encourages interaction" (see Section 7.2.3). The resolution of this apparent contradiction is discussed in the Conclusion chapter to follow.

The completed synoptic diagram was used to present and gain feedback about the study findings in a final visit to the probe participants. Such member checking through graphic elicitation reduces threats to validity arising from researcher bias, since the use of graphic elicitation emphasises and challenges the conceptual structure derived by the researcher. Informants thought that the six token roles (green ellipses) were comprehensive and that their meanings were accurately communicated by diagrammatic context. However, they also noted that *spatial index* and *knowledge handle* could reasonably extend into the social 'sphere', being equally relevant in individual and collective contexts. They thought that engagement was a "sufficiently core concept" to be the central focus of the diagram, and thought that the links between the five most beneficial effects (orange rounded rectangles) reflected the iterative, sequential nature of their approach to teamwork in general. Finally, the four essential qualities of peripheral interaction (red rectangles) were seen to be well motivated by their associated token roles, and an accurate, concise description of peripheral interaction patterns and benefits. Overall, the member checking process served to increase my confidence that the findings as presented in Figure 7.7 were consistent with the experiences and reflections of the users who participated in the technology probe.

7.6 Generality of Results

The foregoing probe analysis, and description of the essential qualities of peripheral interaction, were based on phenomenologically-motivated elicitation of the interaction experiences of a small number of fieldwork participants. The analysis above relates to their experiences alone, and any application of results beyond the particular context studied is contingent on demonstrating a sufficient degree of generality to project these benefits onto other users, interfaces, and contexts.

One way of arguing for such analytic generality is to relate the specific results of a fieldwork deployment to an established theory, thus lending weight to the conclusions drawn. One suitable theory, well established within the field of HCI, is *Activity Theory*. Another suitable theory, taken from cognitive science, is Kirsh's (2001) *The Context of Work*, which explores desk-based office work from a distributed cognition perspective. Both of these theories share conceptual similarities with the account of peripheral interaction as presented here, allowing analytic generalisation to proceed in two orthogonal directions. Firstly, the *essential qualities* of peripheral interaction map onto the interaction model of activity theory, suggesting that these qualities are present to some extent in all physically-mediated interactions. Secondly, the existence and multiplicity of *token roles* in peripheral interaction mirror the use of existing desktop artefacts to provide work context, suggesting that these roles are present to some extent in all desktop-based work contexts. Together, these relationships to established theory suggest that peripheral interaction is not a completely novel paradigm – rather, *it describes a fundamental way in which we already use physical artefacts to mediate our desk-based work activities*. The digital augmentation of physical tokens simply serves to blur the divides between physical and digital resources, internal and external representations, and cognitive and social processes, which constitute the complex, multifaceted nature of working life.

7.6.1 Activity Theory

Activity Theory is a set of conceptual tools, as opposed to a packaged HCI method, which considers motivated actions on objects, mediated by artefacts, as the fundamental unit of analysis. Bertelsen and Bødker (2003) provide a useful overview of the theory in conjunction with a case study of its application. They present a number of characteristic features of activity theory:

1. *Human activity is mediated by socially produced artifacts, such as tools, languages, and representations. This means that, in their immediate relation with their surroundings, human beings extend themselves with artifacts that are both augmentations of and external to the person* (ibid. p305).
2. *Activity is object oriented: It is a (possibly collective) subject's active engagement directed toward an object. This engagement is socially mediated by the community in which the activity is embedded or constituted* (ibid. p303).
3. *Activity is realized through conscious actions directed to relevant goals. Actions are realized through unconscious operations triggered by the structure of the activity and the conditions in the environment* (ibid. p305).
4. *Activity theory understands human beings as dialectically re-creating their own environment. Subjects are not merely choosing from possibilities in the environment, but they are also actively creating the environment through activity* (ibid. p303).

These four characteristics can be seen as corresponding to the four essential qualities of peripheral interaction: FRAME OF REFERENCE (1); FOCUS OF ATTENTION (2); FLOW OF ACTIVITY (3); and FREEDOM OF APPROPRIATION (4) – Figure 7.8 illustrates this as an overly on the synoptic diagram of Figure 7.7. The nature of activity theory as a description of instrument use, either as technical instruments (tools) or psychological instruments (signs), makes it particularly relevant to the use of tangibles since they incorporate tools and signs shaped into the same physical artefact. As such, these qualities, whilst derived from analysis of peripheral interaction experiences, may well manifest themselves in all forms of tangible interaction, albeit to varying degrees. In any case, the fact that they cut across all six of the empirically-grounded token roles, as well as closely resembling the main tenets of an established critical perspective, indicates that the results of the experienced benefits elicitation presented in this chapter are indeed likely to have analytic generality beyond the original fieldwork context.

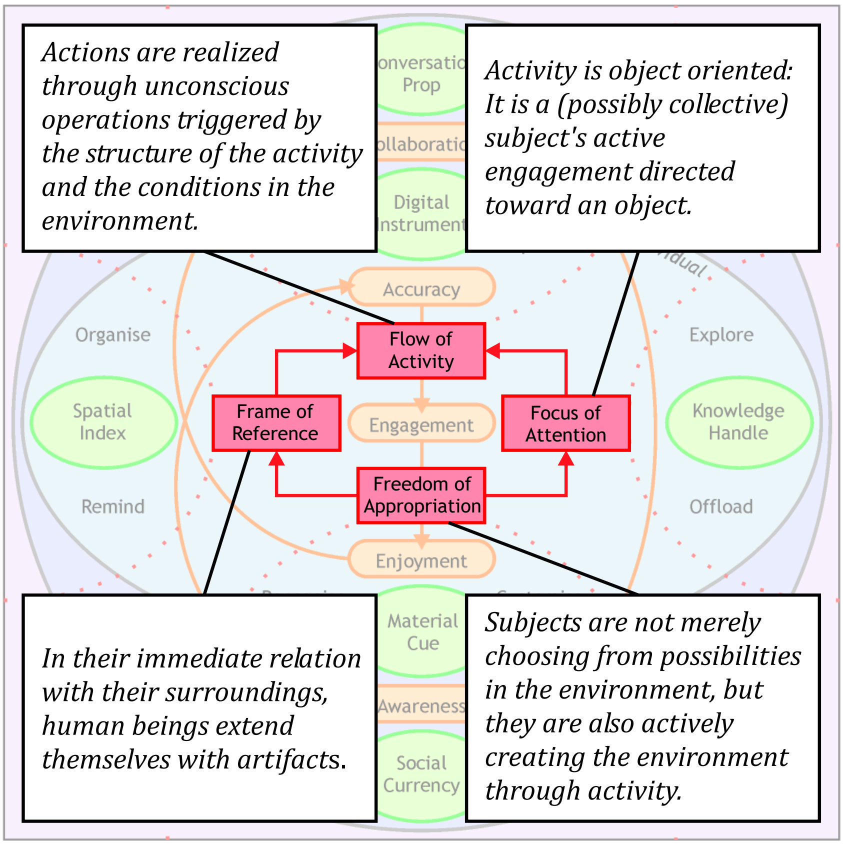

Figure 7.8: Generality of Analysis: Relationship to Activity Theory

7.6.2 The Context of Work

The ubiquitous computing goal of the “portable office”, which can be carried around and used to recreate previous work context through ad-hoc projection onto digitally-enhanced spaces, is predicated on the belief that it is possible to model the deep structure of work context. Kirsh (2001) describes this deep structure as “an underlying system of states, structures, and relations, that can be manifest in offices with different surface features”, and proposes three conceptual devices that abstract away from low-level complexities. These three devices are as follows (ibid.):

1. An *entry point* is a structure or cue that represents an invitation to enter an information space or office task.
2. An *activity landscape* is part mental construct and part physical; it is the space users interactively construct out of the resources they find when trying to accomplish a task.
3. A *coordinating mechanism* is an artefact, such as a schedule or clock, or an environmental structure such as the layout of papers to be signed, which helps a user manage the complexity of his task.

The use of tokens in peripheral interaction is closely related to all three of these conceptual devices. Tokens themselves are a *coordinating mechanism*, with tokens on the interactive surface providing a similar role to the example of a schedule or clock, and with tokens spread around the desktop environment providing a similar role to the example of the layout of papers to be signed. The appropriation of artefacts as coordinating mechanisms is based on their features and affordances; in turn, their appropriation provides *entry points* for future interaction with the coordinating mechanism – Figure 7.9 illustrates this as an overly on the synoptic diagram of Figure 7.7. The user’s current goal within their activity will determine the way in which tokens are coopted as coordinating mechanisms, thus determining which kind of entry point they will provide in future. The six identified roles of tokens in peripheral interaction can be seen as providing different kinds of entry point, corresponding to the underlying kind of coordinating mechanism.

Each way in which tokens can act as *coordinating mechanisms* is related to one of the identified token roles, providing different types of *entry point* for different types of token usage, operating within an *activity landscape* constructed in accordance with the essential qualities of peripheral interaction.

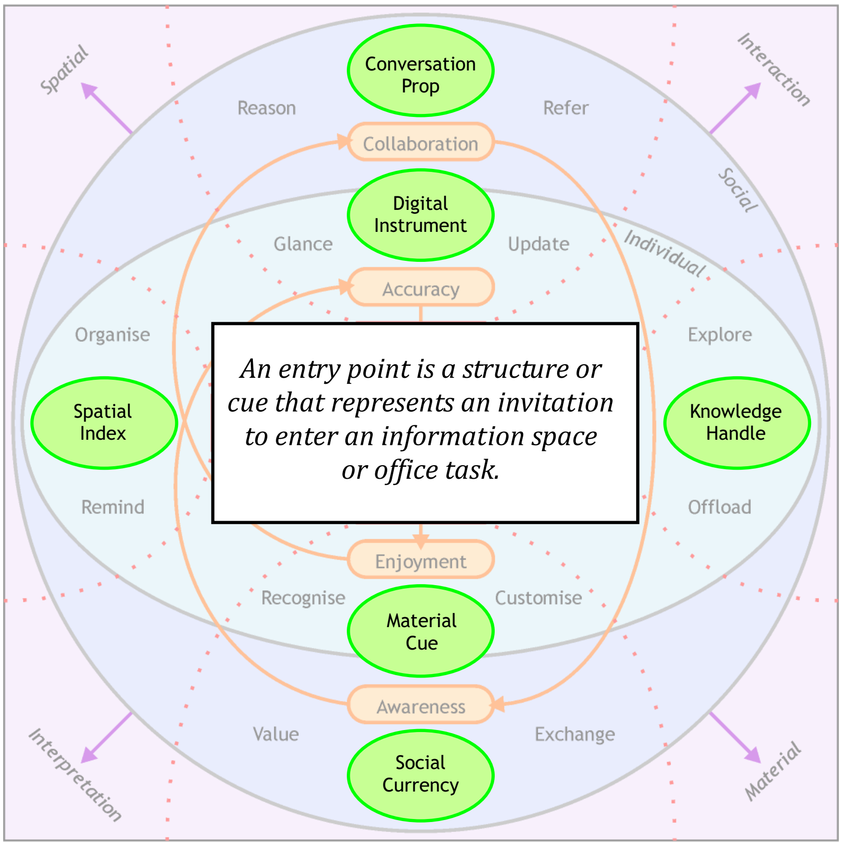

Figure 7.9: Generality of Analysis: Relationship to Entry Points

Chapter 8

Conclusion

This dissertation began with a review of the literature on Tangible User Interfaces, and took a critical view of the research area as a whole. Many published accounts of TUI systems fail to convey the rationale behind the design, or the evaluation of their prototype in action. For TUI research to move forward, it is necessary to begin justifying the benefits of tangibility with respect to real problems in real contexts, and using the experience and insights from doing so as the basis for empirically-grounded theories that future designs can build on.

In my work, I have attempted to make contributions to the design and evaluation of TUIs, through the development of a process for the analytic design of contextually-appropriate interfaces, and a technique for structured reflection on the benefits experienced during interaction in context. However, it is not just the design and evaluation of TUIs that needs addressing, it is the fundamental understanding of what TUIs are, and what they're good for, that should be reviewed.

I believe that two of the foundational characteristics of 'Tangible User Interfaces' need to be re-examined and revised in light of my work. Firstly, I have argued that the original presentation of TUIs as "graspable media", in contrast to "ambient media" (Ishii and Ullmer, 1997), reifies the notions of centre and periphery as fixed categories of the world, rather than treating them as transient states of the mind. Consequently, TUIs are often implicitly regarded as *focal* interfaces that completely engage the user's attention, as evidenced by the overwhelming design of TUIs to support dedicated, focal activities. This is related to the second characteristic of TUIs that I have challenged – the requirement for a *spatial syntax*, in which "the physical state of the interface artefacts partially embodies the digital state of the system" (Ullmer and Ishii, 2001). As well as being focal, TUIs are also implicitly regarded as a way of exploiting *familiarity* with the physical world, by modelling existing physical scenarios or conventions. The domains chosen are predominantly those of an inherently spatial nature, resulting in a spatial syntax that 'naturally' supports *co-located collaboration* around a *shared interface*. Whilst this characterisation of TUIs has served to concentrate research efforts on TUIs that take advantage of the most salient benefits of tangibility, and in a highly successful manner within the HCI community, it has directed attention away from application areas that could potentially make a wider impact.

In this dissertation, I have demonstrated a need for TUIs in the real world of office work, where a focal, shared, familiar spatial syntax for co-located collaboration would be incongruous with the fragmented, concurrent, distributed nature of the activities to be performed. Instead, I have proposed a TUI that is peripheral and personal, with portable tokens that are free from the constraints of a spatial syntax. These tokens can be freely arranged within the sensing region of the interface for ease of interaction, within the boundary of the user's desktop environment for cognitive support, and within their work environment as a whole for social symbolism. The flexibility of such a TUI is necessary to fit in with the complex webs of interactional, cognitive and social activities that make up the multifaceted nature of office work. The design of such a TUI required a contextually-informed, analytic design process and a critical perspective on what it means to be a "tangible" interface.

The theoretical motivation and resulting structure of an asymmetric, bimanual interface for interaction with socially-situated, digitally-augmented physical tokens is a major contribution of this dissertation. I will now present a specific evaluation of the successes and failings of the TUI prototype and its applications (task management, document sharing, and contact awareness), before moving on to give a more general description of peripheral interaction as it was experienced in a real office context.

8.1 The TUI Prototype

Overall, the TUI prototype was a success. However, the relative contributions of the different physical and digital components of the interface to this success vary quite broadly. Perhaps the most significant design decision was to forego a spatial syntax, stepping outside the defined category of "tangible user interfaces" to allow a more relaxed, low-attention means of tangible interaction. This decision gave users the ability to appropriate tokens more freely than what might have otherwise been possible, and the observed instances of users doing so – adapting tokens' spatial and material variables to create meaning beyond the digital – is testament to the value of this freedom. The bimanual style of interaction complementing this lack of spatial syntax, based on a separation of representation (many tokens) and control (single knob), can also be considered a success. Users found the resulting control scheme to be simple, memorable, and efficient, yet safe with respect to only admitting intentional changes. Whilst the "nudge-turn" interaction style is highly specific, the combination of independently meaningful representations and bimanually-safeguarded control is an approach to tangible interaction that could find more general applications.

Reflecting on the usage data for the different token types, task tokens were more successful than document and contact tokens in terms of both the number of tokens in use and the average number of tokens on the surface over time. The likely reason for this is that task token halos contain more information than the halos of the other token types; moreover, this information is of ongoing relevance to each user's current work situation. The ability to see other users' interest in shared documents and the projected workloads of those users are useful features of document and contact tokens respectively, but these features are only of occasional importance when compared to personal task management. The flexible model of task work and means of workload visualisation embodied by task tokens were also seen to be of significant benefit by users. However, it is unclear whether or not an interactive surface was necessary to realise the benefits of the model, compared to the alternative of temporarily co-opting the focal workstation display (or indeed any existing secondary displays). In addition, large-scale project planning and progress tracking may be better suited to large, shared displays than the small personal ones used in the TUI here. Even if changes to the system were made to this effect, tangibles would retain their essential function as a means keeping users aware of their own work plans in a medium that also provides a convenient means of changing them. Such changes may also bring the status of document tokens into parity with task tokens, especially if the document tokens were able to link to existing documents and different document formats.

The primary failing of the TUI prototype was the use of token edge texture to identify the owners of task tokens and the people represented by contact tokens. Aside from the difficulty in scaling such an approach beyond the small number of edge textures that are easily distinguishable from a cursory glance, the social implications of making concrete and undeniable the sometimes vague or undefined notion of task 'ownership' were not fully considered during the design process. Conversely, users found the token edge textures to be both visually aesthetic and haptically pleasing, therefore they may well have been better used to represent known, fixed categories of each token type, such as task importance (high, medium, low) and document permission (just me, my work group, any work group). Aside from other minor failings that could easily be corrected in future prototypes – the inability to link to any document, the lack of integration with work calendars and task management software, and an inexpressive calendar tool halo – the system largely supported the style of interaction for which it was intended.

8.2 Peripheral Interaction

My conceptualisation of peripheral interaction began as a description of the fast, frequent nature of interactions required in the safety-critical context of Air Traffic Control. The introduction of peripheral interaction to the less-critical office context was justified by its use in performing auxiliary work activities in parallel with primary tasks – in such situations, I thought interaction efficiency would be key in encouraging frequent switching between interfaces, furthering the appreciation and adoption of the TUI. However, after an extended deployment of my prototype TUI in the context of a real office, probe participants did not rate efficiency as an important factor in contributing towards the benefits of the interface. Whilst fast, frequent interactions are indeed possible with the TUI, participants did not see this aspect of the interface as being the primary justification for its existence in the workplace.

In addition to periodic, low-attention interaction with digital information, users can engage tokens for the material, spatial and social interpretations they have given them. This can be by individuals alone at their desks, as well as by groups of individuals in meetings and other encounters. All uses of tokens are predicated on the freedom with which users can adopt and adapt physical tokens for their own needs, appropriating tokens' graphical, material, and spatial variables to externally represent structures of cognitive and social significance.

Interaction with such structures takes place on an *engagement spectrum*. At the highest level of involvement, tokens can provide a focus of attention for users through manual handling, visual fixation, and verbal deixis. Users can engage with tokens for their physical properties, digital content, mnemonic function, or social connotations; they can act as both reminders for future actions, and assist recall of prior interaction history.

At the opposite end of the engagement spectrum there are tokens providing a tangible frame of reference for work activity, selectively reflecting the most important aspects of the current work situation. This salient external structure resides for the most part on the periphery of users' workspace and attention, ready to lower the barriers to potential future interactions. These interactions are not just with digital information, but with other people about the things represented by the tokens.

The reduced costs of interaction arising from an external frame of reference, combined with its ability to focus users' attention, serve to ease the flow of thought, action, and information between users and their TUIs, between users and their desktops, and between users themselves. Tokens, and the things they represent, can drift in and out of users' attention in synchrony with the ebb and flow of their work activities. Peripheral interaction is not just about low-attention interaction with digital information, but about being able to selectively and fluidly engage tokens on multiple levels of meaning. It is these qualities that "sustain users' attention and encourage interaction". This engagement is not based on extended periods of concentration, but on the constant, peripheral availability of tokens for immediate yet fleeting bouts of attentional focus. Moreover, this peripherality does not appear to be confined to interaction with my TUI; rather, it can be seen as a reflection of the way in which we engage with most things in the real, physical world. Peripherality is the norm, rather than the exception, but the absorbing nature of traditional interactions with the digital world obscures this reality. I argue that the qualities of peripheral interaction presented here are in fact the fundamental qualities of tangibility that distinguish it from virtuality, and that no amount of ubiquitous, multi-touch display surfaces will provide the same degree of harmony between users and the information in their environment. To quote a participant in my study, "even if the screen was perfect, at a glance I still wouldn't see it as easily as something physical".

8.3 Addressing the Research Questions

The original goals of this research were to answer the following research questions:

1. APPLICATION CONTEXT. What contexts, in terms of activity and environmental structures, are best suited to peripheral interaction?
2. INTERACTION STYLE. What procedures, tools and methods should we use to design for a peripheral interaction style in these contexts?
3. INTERFACE STRUCTURE. What structural forms, styles of mapping, and modes of representation can support this peripheral interaction style?

4. INFORMATION CONTENT. What are the different kinds of information that this peripheral interface structure can represent?
5. APPLICATION JUSTIFICATION. What are the essential qualities of peripheral interaction that justify its use in interacting with such information?

Whilst this dissertation has provided substantial answers to each of these questions, naturally there are areas that could be investigated further.

Regarding Research Question (1), I have given a thorough discussion of why peripheral interaction is a valuable capability in the office context. In the final section on future directions, I will briefly list a number of alternative contexts in which peripheral interaction may be beneficial, although these ideas are speculative have not been empirically evaluated.

For Research Question (2), I have drawn together many theories from the area of TUIs, from HCI in general, and from a wider range of traditional disciplines. I have demonstrated the utility of such a process by applying it to the context of desk-based office work, and using the component elements of the process to analytically support the resulting design of a TUI for peripheral interaction. However, the process is yet to be tested in contexts other than the office.

My asymmetric, bimanual TUI structure has been shown to support peripheral interaction in real contexts, and this partially addresses Research Question (3). However, it has also been shown that the means of tangible interaction is potentially less significant than the simple existence of digital and social information in physical form. There may be other interface structures that adequately support a peripheral style of tangible interaction, both on the conventional office desktop and beyond.

To answer Research Question (4), I have attempted to identify the different ways in which users can incorporate tokens into their working life; these "token roles" reflect the way in which tokens can be seen as conveying different kinds of information in different interactional situations. These token roles are grounded in the coding of interview transcripts from the field deployment of the TUI, but deployment in other contexts may have resulted in different interpretations. However, the way in which I was able to organise these roles into a coherent diagrammatic representation suggests that they are fundamental categories of token use, instances of which were noted and reflected on by fieldwork participants.

Finally, I believe that Research Question (5) is addressed by the "essential qualities" of peripheral interaction, which provide both justification for its existence as a recognised style of interaction, and rationale for its use in contexts where those qualities are appropriate. Perhaps the greatest current limitation of this research is the fact that the extent of the evaluation was known in advance to be of fixed duration, and that users of the TUI always knew they would have to give it up after the study had reached its conclusion. Whilst the log data was detailed, the interview responses rich and insightful, and the information managed of real importance, personal investment in the TUI never quite reached the level of a technology to be used on an ongoing basis. Multiple long-term studies of peripheral TUI usage would be necessary to investigate whether the identified "essential qualities" of peripheral interaction are experienced more generally.

8.4 Implications for Design

The four "essential qualities" of peripheral interaction derived from analysis of the probe and elicitation results naturally translate into concise implications for design:

1. *Tangibility is about giving users the freedom to project new systems of meaning onto the world, based on the existing meaning that it already has for them*. The ways in which users can do this – by appropriating the graphical, material, spatial and social variables of tangible objects – are limited by the constraints of spatial syntax. Designers should aim to strike a balance between functionality and freedom: whilst functionality encourages initial adoption of a technology, freedom of appropriation is what determines long-term investment and meaningful use.

2. *Tangibles are not just physical objects with digital augmentations, they are socially-situated objects of external and distributed cognition*. When a user focuses their attention on a tangible object, they may be considering its material construction or its spatial location on either a literal, indexical or symbolic level, as something that stands for itself, for its digital context, or for its cognitive or social function. Designers should consider the trade-offs between all of these levels and types of meaning.

3. *Tangibles exist in the physical world and have the potential to remain meaningful even when outside of the sensing region of their parent interface.* This blurs the distinction between interface and environment, and allows tangibles to passively yet persistently frame the context of work. Such framing operates on multiple levels: it can lower barriers to future interactions; it can provide an external memory of work progress and plans; and it can reflect the standing of an individual within a social group. Designers should consider the ways in which tangibles can decorate their context of use in ways that are peripheral yet supportive to the main activity.

4. *Tangibility provides an opportunity for selective, fluid, episodic engagement with information in a more direct manner than WIMP-based interaction*. The structure of users' local environment can be set up to support the structure of activities, guiding the user and guarding against the effects of interruption and memory failure. Tangibility also provides an opportunity to engage with other people on a social level – initiated, mediated and recorded by the presentation, discussion and exchange of tangible objects. Designers should consider how the performance of work in real contexts is influenced by environmental factors, and how the use of tangible objects can help regain and retain the flow of activity.

The arrangement of "token roles" into a synoptic diagram (see Figure 7.7) also leads to an associated implication for design – that *designers should consider where they intend their interface to lie within the abstract "usage space" defined by the dimensions of spatial-material, interaction-interpretation, and individual-social, what the effects are of design changes that move the interface into new areas of this space, and what the resulting balance is between appropriately supporting user activities and supporting the activity of user appropriation*. Returning to my TUI prototype, the initial conception of the system was as a way of supporting auxiliary work activities using tangibles. Whilst its evaluation showed the interface to be appropriate in this respect, it also showed that the interface encouraged appropriations that extended it into all areas of this abstract usage space. If I had designed a peripheral GUI, or indeed a conventional TUI that tied the tokens into a spatial syntax, the interface would not have supported the diverse range of user appropriations, in terms of token roles, that ultimately resulted. However, the potential for inadvertently supporting inappropriate user activities is also something that should be considered, underlining the value of systematically analysing TUI designs from all areas of this space.

The particular token role of social currency also hints at a more fundamental difference between graphical and tangible interfaces: tangibles have symbolic value arising from their inherent costliness. This cost arises from multiple sources, including the economic cost of the materials used to construct the tangibles, the opportunity cost of the space that they occupy, and the opportunity cost of the time used to set up their physical-digital mappings. It is these costs that mean tangibility "doesn't scale", i.e. it is infeasible to have a physical object for every potential digital referent. However, even if tangibility could scale in such a way, creating too many physical-digital mappings would diminish their symbolic potency. As it stands, in the words of P3, "You have to choose what's most important at that point in time". On reflection, my TUI embodies multiple levels and types of importance. In terms of levels of importance, the token types are the general classes of thing that are of ongoing, shared importance in the context of office work, whilst in any given office, the set of instantiated tokens represents a snapshot of the specific things that are currently of most importance to the people who created them. In terms of types of importance, token roles are the different kinds of engagement with instantiated tokens from which users derive value. The cost of tangibility should be matched against things of importance to the user, and the value derived from instantiating a tangible must outweigh the costs of doing so. The final implication for design, therefore, is that *tangibility has symbolic value that is difficult to replicate digitally – this value should be used to support more meaningful, calm engagement with the things of importance to users in their spatial and social contexts.*

8.5 Future Directions

Future work on peripheral interaction needs to investigate long-term usage of peripheral TUIs in context, as well as the network effects of greater deployment coverage within organisations. Applications that take advantage of loosely coupled 'tangible workgroups' to exploit latent organisational knowledge via peripheral communication channels would be a natural extension of socially-motivated peripheral interaction.

Peripheral interaction in contexts other than the generic office is also worthy of investigation. The current TUI applications could be adapted to support work in specific domains – for example, a solicitor's TUI would emphasise tangibility as a means of direct, physical time tracking for billing purposes and client file access. The TUI structure could also be modified to suit work in contexts that are not desk-based, such as order tracking in restaurant kitchens. As well as work contexts, peripheral interaction could be used in the home, providing tangible access to media, recipes, and communication with friends and family. The use of peripheral interaction in computer gaming, as a strategic overview to complement the immersive nature of regular tactical gameplay, is also an interesting opportunity for further work. Whether in the home, office, or elsewhere, a prominent design objective should be to create compelling social experiences through interaction with tangibles, using their ability to bring people together in joint engagement. Even for the predominantly lone solicitor, chef or gamer, the physicality of tangibles means that they can act as a 'social currency' that people value jointly for their physical form, digital content, and their symbolic interpretation within a social group, ultimately providing additional opportunities for face-to-face interaction.

One of the two biggest challenges for TUI research is the integration of interactive technologies – getting the personal and shared, focal and peripheral, graspable and ambient, and local and remote interfaces to support the distributed solution of problems beyond the capabilities of any individual person or technology. This is essentially the superordinate goal of ubiquitous computing, and whilst technological advances will make such integration possible, contributions from HCI will be necessary in order to render the resulting 'solution' appropriate to not just the problem, but the problem context. The second major challenge, for TUIs in particular, is the identification and exploitation of uniquely physical phenomena in a manner that justifies the use of tangible interfaces and differentiates them from their graphical counterparts. In this dissertation, I have identified and supported a peripheral style of interaction that exploits some of the unique characteristics of tangible objects, but substantial progress is needed before we can say that we truly understand the benefits of tangibility.

Appendix A

Supplementary Material

A.1 Tangible Correlates Analysis

Sample Tangible Correlates analysis, taken from Edge and Blackwell (2006).

A.1.1 Continuous values: spatial approaches

Continuous attribute values associated with an information entity can be represented either by position and orientation of objects within some spatial frame of reference, or as mechanical properties of the physical objects used. The advantage of the spatial alternative is the high $\texttt{visibility}_{\texttt{<CD>}}$ and $\texttt{juxtaposability}_{\texttt{<CD>}}$ that result, and the potentially high $\texttt{structural correspondence}^{[\texttt{closeness of mapping}]}_{\texttt{<TC>}}$ from representing continuous values – if they represent spatial quantities – in a spatial manner. However, two entities may share the same value for a particular attribute, but two physical objects cannot share the same position in three-dimensional space. This means that $\texttt{structural correspondence}^{[\texttt{closeness of mapping}]}_{\texttt{<TC>}}$ when representing non-spatial properties (or spatial properties at too low a resolution) may be adversely affected.

A further advantage of the spatial approach is the ease with which three dimensions of information can be set simultaneously – i.e. the low degree of $\texttt{rigidity}^{[\texttt{viscosity}]}_{\texttt{<TC>}}$. The tradeoff is that it may be difficult to adjust any one of these dimensions independently – relative to a fixed surface and user position, any gross movement of the arm will necessarily produce a rotation of the hand, which must be countered by fine movements of the fingers – a manifestation of $\texttt{unwieldy operations}^{[\texttt{hard mental operations}]}_{\texttt{<TC>}}$ resulting in $\texttt{shakiness}^{[\texttt{error proneness}]}_{\texttt{<TC>}}$. This has other implications too – if at most three dimensions of information can be set in this way, but more than three dimensions need to be represented, the others will need to be represented in some other manner, reducing $\texttt{consistency}_{\texttt{<CD>}}$. This is due to the limited availability of spatial dimensions as a representational medium.

A.1.2 Continuous values: mechanical approaches

We are particularly interested in opportunities to represent continuous values by changing the physical configuration of an object. Physical configurations can be decomposed into fundamental kinematic pairs of elements (Reuleaux, 1876). The two parts of a kinematic pair are constrained in the way they move relative to each other, and the resulting degrees of freedom can be used to represent one or more continuous values. Consider a single value represented by a screw pair, such as a nut and bolt, whose relative motion describes a continuous helix. If we imagine the pitch of the screw threads getting increasingly shallower, the nut will eventually turn without moving along the screw, becoming a rotational pair. Alternatively, imagine the threads stretching out until they become grooves along the length of the screw, in which case the nut would slide along without turning, forming a translation or prismatic pair. Hence, screw pairs, revolute pairs and prismatic pairs are the fundamental constructions for physically representing onedimensional values in tangible interfaces, with twisting, turning and sliding being the fundamental actions for manipulating them.

We can make a CDs comparison between these three one-dimensional kinematic pairs by considering simple expressions of each. A sliding pair can involve one part sliding within another – a "position-slider" – or two parts sliding relative to one another, making a "length-slider" (or "telescopic") device such as the stretchable square in the Bricks TUI made by Fitzmaurice, Ishii & Buxton (1995). A turning pair can be a direct rotation device – as in a "knob" – or an indirect rotation device, where rotation of the "joint" is a consequence of the movement of the joined elements (an "articulated" tangible). In an analogy to the linear arrangement of sliders, "position-screws" represent some quantity by the position of a nut on a bolt, and "length-screws" by the extension of a bolt out of a threaded cavity (like a swivel chair which uses a screw pair for height adjustment). Other simple expressions of the basic pairs are possible, but we will focus on these six. The most fundamental trade-off is between the $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$ of the physical representation – the amount of information expressed per unit of length or space – and its $\mathtt{rigidity}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$, or how much it resists change. The compact linear form of a position-slider means that multiple sliders can be placed side-by-side, for simultaneous operation with simple hand movements. The interface of audio control devices such as graphic equalisers and mixing desks is a good example. In contrast, a length-slider has more $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$ due to its varying size, and takes more time to operate if it requires two hands, increasing its representational $\mathtt{rigidity}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$. A knob can have even less $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$ than a slider, but takes slightly more time to operate since there are fewer tactile cues as to the current value – a marginal increase in $\mathtt{rigidity}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$. A joint requires two hands to operate and so has a similar rigidity to a length-slider, but takes up varying amount of area depending on the joint angle, increasing its relative $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$. Screw pairs generally have less inherent $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$ than sliders due to the use of an extra dimension when expressing the single degree of freedom (think of a screw pair as a coiled-up slider). However, they have more inherent $\mathtt{rigidity}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$ than the other pairs, because many rotations may be required to achieve a given translation.

The situation can be visualised as shown in Figure A.1.

There are many other tradeoffs associated with these kinematic pairs. Joints can be composed into a linkage, providing potential for $\mathtt{adaptability}^{[\mathtt{abstraction}]}_{<\mathrm{TC}>}$. Length-sliders and length-screws have the greatest degree of $\mathtt{role\ expressiveness}_{<\mathrm{CD}>}$, as they represent quantity by changes in physical size. Screw pairs are difficult to change quickly, and so display the least $\mathtt{shakiness}^{[\mathtt{error\ proneness}]}_{<\mathrm{TC}>}$. Position-sliders and position-screws have a greater degree of $\mathtt{juxtaposability}_{<\mathrm{CD}>}$ when arranged for side-by-side comparison. Pairs in which only one part touches a surface (position-sliders, knobs and some position-screws) have less $\mathtt{rootedness}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$ than those in which both parts rest on the surface, because movement is less likely to affect their configuration. Knobs have relatively low $\mathtt{rigidity}^{[\mathtt{viscosity}]}_{<\mathrm{TC}>}$ and $\mathtt{bulkiness}^{[\mathtt{diffuseness}]}_{<\mathrm{TC}>}$ and can also be tailored in different ways. Although a simple knob can only express values within the range of a single revolution (about as expressive as a

Continuous attribute conditions	**Recommended syntax**
Low space constraints	Spatial position
One or two attributes to express	
Tokens unlikely to share the same values	
One attribute to express	Orientation
Little movement of tokens necessary	
Side-by-side attribute comparison	Position sliders
Rapid operation	
Peripheral comprehension	Length sliders
Rapid operation	
Side-by-side attribute comparison	Position screws
Accurate operation	
Peripheral comprehension	Length screws
Accurate operation	
Multiple related controls (coaxial or linear)	Knobs
Rapid operation	
Multiple related controls (coaxial or linear)	Joints
Peripheral comprehension	

Table A.1: Summary of Syntax Recommended for Continuous Attributes

position-slider), knobs can also be augmented to track the number of revolutions, for example with an array of lights. This simulates the information range of a screw pair. A more abstract virtual layer might allow the knob to exploit its free-turning property, allowing an unbounded range to be represented. However, getting to a value outside the expected range might be a timeconsuming if the angular increment of the knob is inappropriate. In this case we might use coupled combinations of knobs, for example controlling logarithmic increments of 1000, 100, 10, and 1. This may result in faster and more accurate control, but the physical state no longer corresponds directly to the value controlled, introducing another system of tradeoffs between $\texttt{rigidity}^{[\texttt{viscosity}]}_{<\texttt{TC}>}$, $\texttt{adaptability}^{[\texttt{abstraction}]}_{<\texttt{TC}>}$ and $\texttt{structural correspondence}^{[\texttt{closeness of mapping}]}_{<\texttt{TC}>}$. Alternatively, one knob might define a multiplier ratio for the other (like the front and rear gears on a bicycle), as in the SeismoSpin device (McKelvin et al., 2003) designed to navigate time on a scale of minutes to decades (Table A.1).

A.2 Embodied Nature of Peripheral Interaction

The bimanual scheme of interaction can also be explained in terms of image schemata and metaphorical mappings, even though it was not conceived with these in mind. Using the image schemas of Lakoff (1987):

1. *Physical tokens have digital attributes*. Token attribute halos employ a CENTRE-PERIPHERY schema: the central physical tokens are more important than their surrounding digital attributes. This schema derives from our bodily experience of having trunks that are more important for identity and survival than our limbs.
2. *Token nudge is attribute selection*. Attribute selection employs a SOURCE-PATH-DESTINATION schema and the PURPOSES ARE DESTINATIONS metaphor. Our bodily experience of moving ourselves and moving other things to DESTINATIONS in the world is almost always motivated by a PURPOSE, and can be directly understood without a metaphorical mapping (ibid.). Nudging is a kind of movement, and SELECTION is a kind of purpose, so the process of nudging to select an attribute can be understood metonymically as MOVEMENT from the SOURCE (by NUDGING the token), along a PATH (in the attribute's DIRECTION), to the DESTINATION (the PURPOSE of attribute SELECTION).
3. *Knob rotation is attribute manipulation*. Control of the selected attribute is mapped to the single physical knob. The resulting bimanual interaction employs a LINK schema connecting attribute selection with one hand, via our body, to attribute manipulation with the other. This schema is bodily experienced in early life through the umbilical cord link and the holding onto parents and things to secure our connection to them. The control schema also employs a CLOCKWISE IS MORE metaphor for manipulating attribute values.

A.3 Calendar Pseudo Code

```
calculateLatestTaskRestartTimes()
  order tasks from last to first due
  dayPointer := latest due date of all tasks
  taskIndex := 0
  while (taskIndex < size(tasks) and
         tasks[taskIndex] is due on dayPointer)
    if (task has zero duration)
      setStartDate(task, getDueDate(task))
    else
      add task to activeTasks[dayPointer]
      taskIndex := taskIndex + 1
    minsPerTaskThatDay := 0;
      if (at least one active task)
        minsPerTaskThatDay := workingMinutesEachDay[dayPointer] /
                              size(activeTasks[dayPointer]
        while (true)
          find all active tasks that finish on dayPointer
          if (at least one finisher)
            for (each finisher in order of first to finish)
              extraMins :=  minsPerTaskThatDay -
                            minsRemaining[finisher]
              remove finisher from activeTasks
              minsPerTaskThatDay := minsPerTaskThatDay +
                                    extraMins/size(activeTasks)
              setStartDate(finisher, dayPointer -
                                     (minsRemaining[finisher] *
                                     (size(activeTasks)+1))
              minsRemaining[finisher] := 0
            continue with remainder of this day
          else
            for (each active task)
              minsRemaining[active] := minsRemaining[active] -
                                       minsPerTaskThatDay
            break to next day
        if (no active tasks and dayPointer < earliest due date)
          return
        else dayPointer := dayPointer - 1 day
```

A.4 Contact Tokens and Task Delegation

Figure A.1 shows how task and contact tokens work together to help coordinate work efforts. Photograph (a) shows the surface of a user called Laura inspecting the timelines of two other users named Steve and Angela, which are displayed above Laura's own timeline along the top edge of her interactive surface. The timeline of the last nudged contact token has its background shaded with light grey so that users can easily distinguish between multiple timelines, and associate them with their corresponding contact through the contact token. In the case of photograph (a), the contact token for Steve was last nudged, highlighting his two upcoming tasks of "Debugging" and "Documentation".

Photograph (b) shows the effect of Steve delegating both of those tasks to Laura, which is inferred from their appearance on Laura's interactive surface. The task tokens clearly belong to Steve since they share the same edge texture, and they are integrated into Laura's timeline according to her specified working hours. Tasks delegated to a user are shown as dark green boxes in the user's timeline, to differentiate them from the grey boxes of regular tasks when the related task tokens are not on the surface. The detection of the "Debugging" and "Documentation" task tokens on Laura's surface will trigger a sequence of updates in which the delegation is propagated to all interested parties, who will observe the transfer of the task from Steve's timeline to Laura's timeline as soon as Steve's interactive surface has recalculated the timeline visualisation and uploaded the changes to the server. The consequence of this is that shortly after photograph (b), the "Debugging" and "Documentation" tasks still visible in Steve's calendar will disappear, indicating to Steve that the task tokens have in fact been registered on Laura's interactive surface. The converse of this is that if Steve's surface is not logged into the server, it will not be able to recalculate his calendar and Laura's surface will remain as in photograph (b). The persistence of duplicate task entries in this case is an indication to Laura that Steve has not been made aware, via his interactive surface, that the tasks have been successfully integrated by Laura into her own timeline.

Photograph (c) is of Steve's interactive surface, who is currently observing Laura's timeline visualisation of workload. The two tasks of "Debugging" and "Documentation" are shown in dark green to indicate the fact that they are tasks delegated by Steve to Laura, mirroring their same green appearance in Laura's timeline. Finally, photograph (d) shows the effect of Laura handing these tasks back to Steve, at which point they lose their delegated status (since they are his task tokens) and are displayed in the workload visualisation as regular grey boxes. Photographs (e) and (f) provide an example of the new forms of social interaction facilitated by physical task representations: the returning of delegated task tokens by playfully throwing them to a team member across the room.

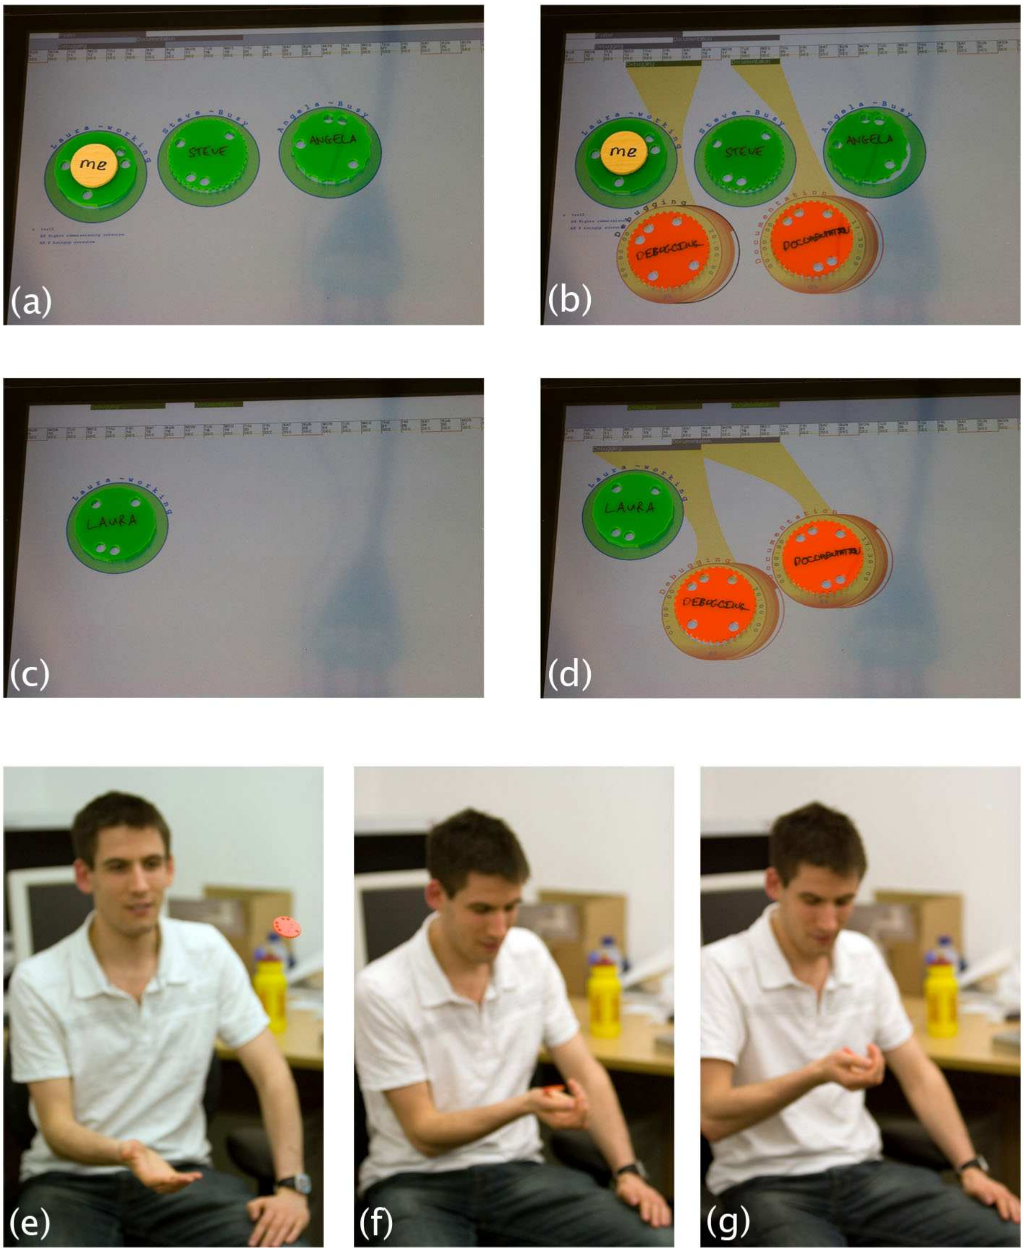

Figure A.1: Task Delegation and Tracking using Task and Contact Tokens

A.5 Calendar Tool

Interaction with the calendar tool is similar to interaction with task tokens. Nudging the calendar tool upwards displays a link to the currently selected day, which can be moved forwards and backwards in time by rotating the Powermate as shown in photographs (a,b,c,d) of Figure A.2 (the red calendar tool is augmented with a black material identifier). Once the desired day is selected, the Powermate can be used to adjust the anticipated number of working hours – that is, the number of hours on that day which will be dedicated to the completion of tasks represented in the system. A demonstration of the adjustment of working hours on a single day is shown in photographs (e,f,g,h).

The currently specified number of working hours on the selected day is shown along the top edge of the calendar tool, and indicated more permanently by the vertical displacement of a line running horizontally through each day of the timeline. When this line is yellow, it means the number of working hours on that day follows the default value. By performing a press-turn action with the Powermate, in which the Powermate is rotated whilst pressed down, the number of specified working hours on the selected day can be manipulated. In this situation, the line-based indication of working hours for that day turns red. A simple click of the Powermate on any selected day will cause it to revert back to yellow and to the default number of working hours. Photographs (e,f,g,h) in Figure A.2 shows the effect of such an adjustment on the task named "Talk" – as the anticipated number of working hours on its planned completion date is reduced from seven to one, its latest restart date moves back in time, closer to the present. Notice the large jump between photographs (g) and (h), caused by the default of zero hours work planned for days of the weekend – such non-linear behaviour is difficult to simulate mentally, yet can be freely explored through the interface.

The dynamic visualisation of non-linear behaviour is true to an even greater extent when the default number working hours per day is manipulated, as shown in photographs (i,j,k,l) of Figure A.2. Here, the calendar tool has been nudged to the right, towards a vertical-bar representing the default number of working hours per day. Subsequent rotation of the Powermate has the local effect of adjusting the height of that bar according to the new default number of working hours, and the global effect of adjusting the vertical displacement of all the yellow lines running through the timeline, which reflect this default value. Experimenting with different numbers of default working hours per day can be seen as a form of sensitivity analysis that allows users to visualise the possibly dramatic effects of seemingly small changes in planned work output.

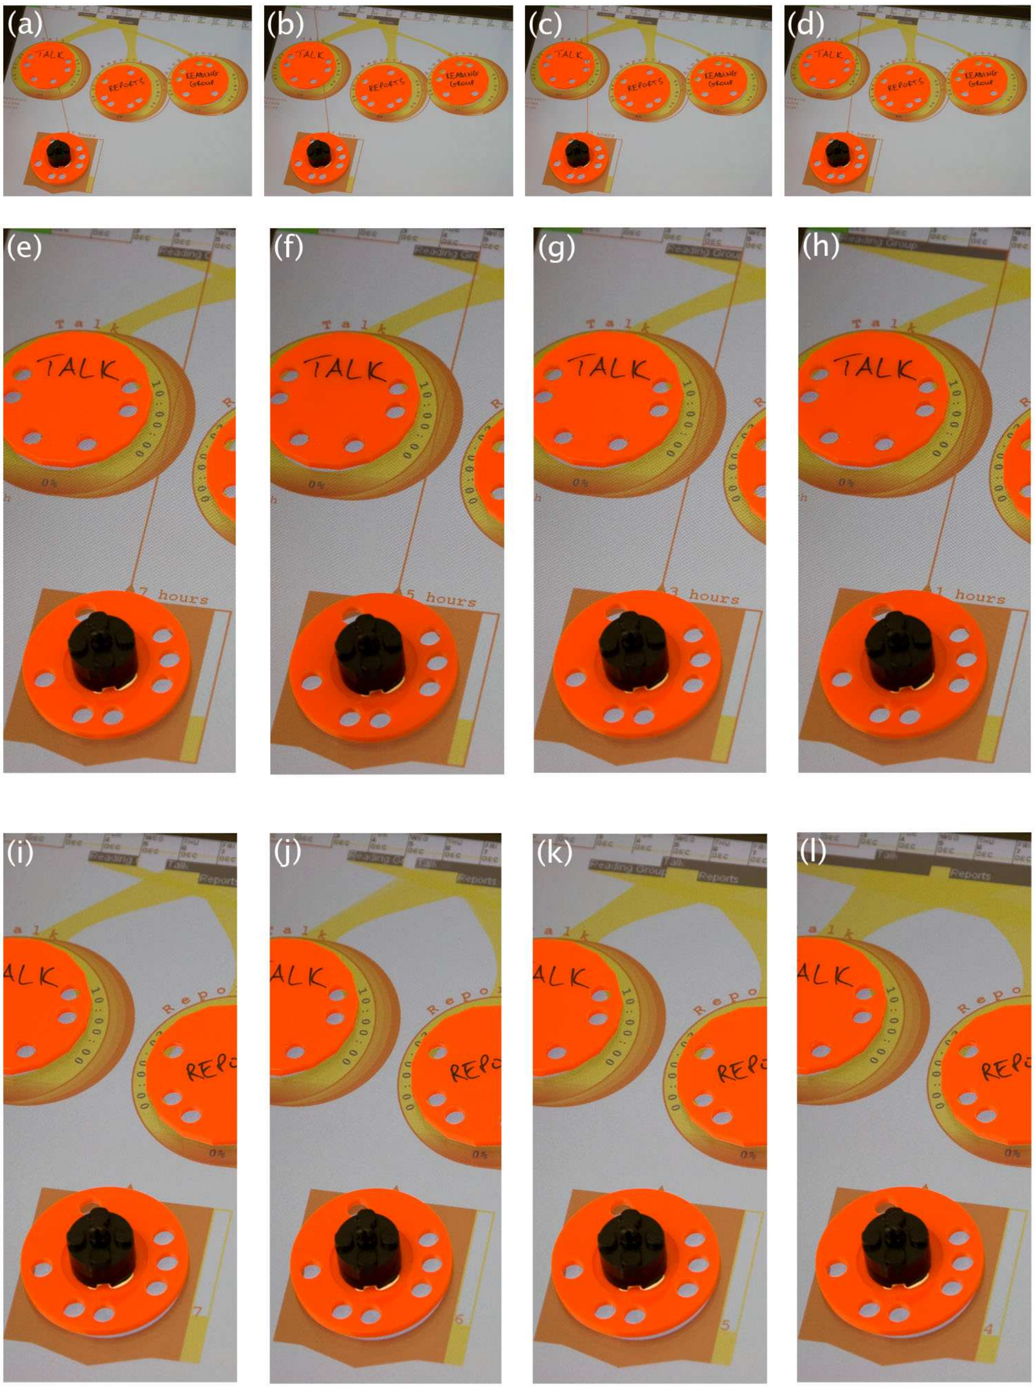

Figure A.2: Adjustment of Hours per Day using Calendar Tool

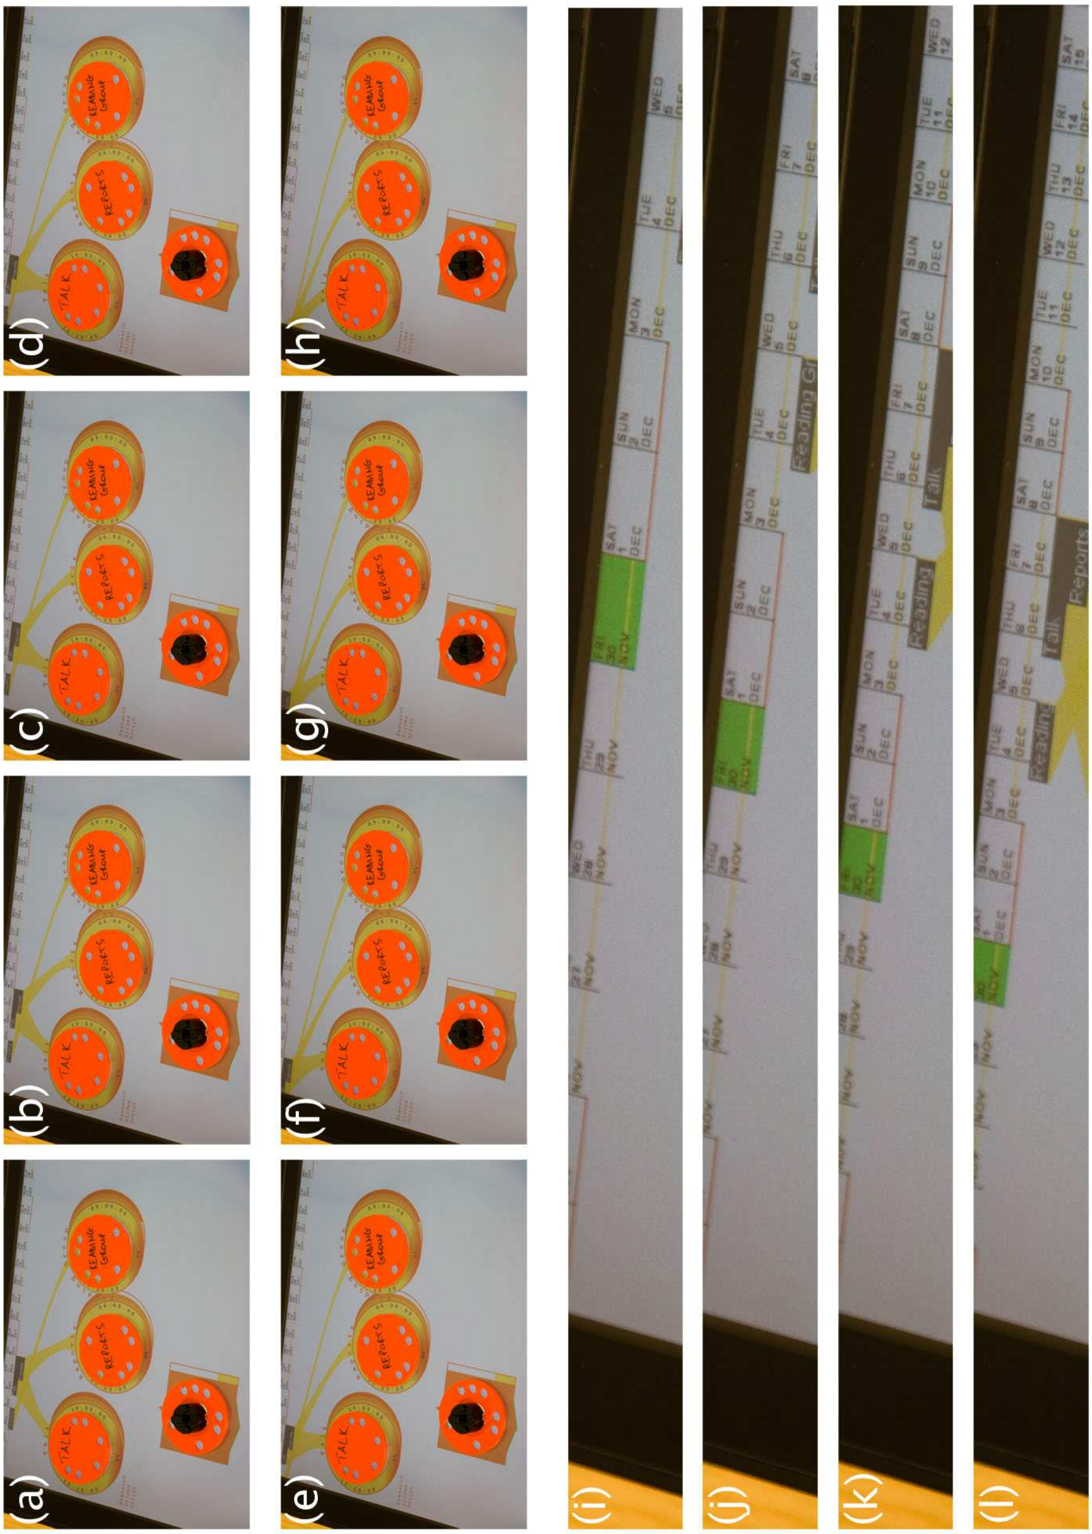

Figure A.3: Timeline Navigation and Scaling using Calendar Tool

As well as allowing users to adjust their planned working hours, the calendar tool also allows users to navigate and zoom the scale of their timeline. These functions are selected by nudging the calendar tool downwards: subsequent rotations of the Powermate scroll the timeline the navigation purposes (Figure A.3, top) whilst press-turn actions adjust the size at which individual days are displayed in the timeline (Figure A.3, bottom).

Within the timeline, the current day is indicated by green shading between the lines that delimit the start and end of the day. The default presentation of the timeline is such that the current day is at the top left-hand corner of interactive surface, allowing users to visualise their workload a customisable number of weeks into the future. A short click on the Powermate after it has been nudged downwards returns the timeline to this default presentation.

A.6 Computer-Vision Management of Tokens

In this appendex, I describe the underlying computer-vision algorithms for token management, ranging from the initial detection of tokens, to their subsequent identification and tracking. My goal was to create a functioning prototype that was robust enough to withstand field deployment – the algorithms presented here therefore represent what was necessary, rather than what could potentially be achieved using state-of-the-art techniques.

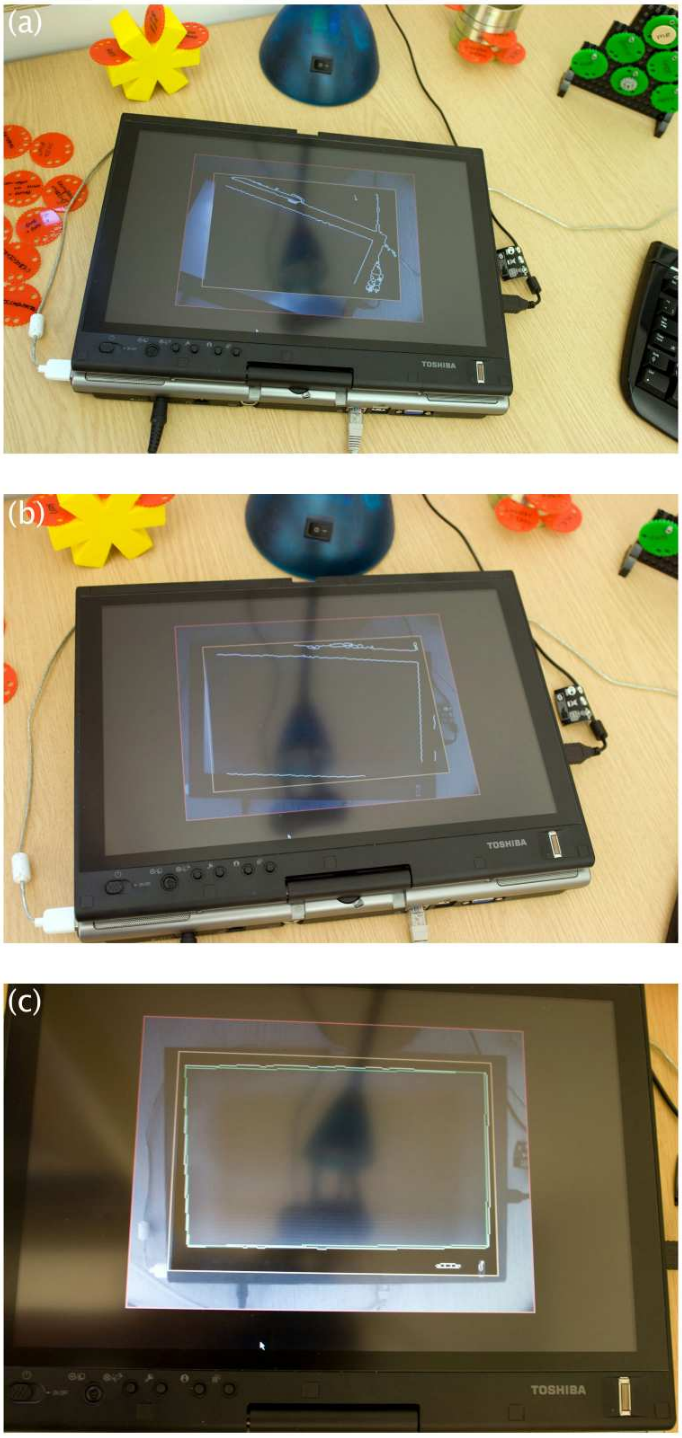

Figure A.4: Interactive Camera–Surface Calibration: (a,b) feedback as to the relative positioning of the Webcam and tablet PC guide the user to slide the tablet PC into the correct alignment; (c) all four corners of the tablet PC are within the target region – the user may proceed to system usage.

A.6.1 Camera–Surface Calibration

Before any token-based interactions can occur, the spatial relationship between the Webcam and the tablet PC needs to be established. Just as the detection and identification of tokens makes use of the characteristic, brightly-illuminated background of the interactive surface, so too does the calibration of the camera relative to that interactive surface. To achieve a camera position capable of pointing directly at the centre of the interactive surface, the Webcam is anchored to an unlit desk lamp flexed above the desired location of the surface. Following initialisation of the interactive surface application on the tablet PC, the exact positioning of the tablet PC beneath the Webcam is guided by the following interactive process:

1. The upwards facing screen of the tablet PC flashes white for a fraction of a second. During this time, the Webcam takes a snapshot of the tablet PC's current position and orientation.
2. The high degree of contrast between the bright white screen of the tablet PC and the black border surrounding it presents a simple target for edge detection[1].
3. Feedback from the image processing, in terms of the processed Webcam image, is presented to the user for a couple of seconds before the next screen flash and image capture.
4. When a bright rectangle is detected such that the four corners of the rectangle fall within some threshold distance of the four corners of the Webcam image, calibration has been achieved and usage of the system may proceed.
5. The user can fine tune the position and orientation of the tablet PC until they are satisfied with the alignment, at which point they can switch into token interaction mode using the miniature joystick on the casing of the tablet.

This process is illustrated in Figure A.4. The outer red rectangle is the boundary of the Webcam image as it is presented on the interactive surface, necessary because the Webcam and tablet PC do not share the same aspect ratio. The inner amber rectangle represents the target area for the screen of the tablet PC, determined by the effective cropping that results from the smoothing of image brightness. A green quadrilateral linking the four detected corners of the tablet PC complete the Red–Amber–Green traffic-light metaphor, signifying that the calibration is complete and that the user is free to proceed into token interaction mode.

The Java Advanced Imaging (JAI) API is used to calculate the required perspective correction based on the deformation of the tablet PC as it appears in the Webcam image, which transforms coordinates from the Webcam image into the rectangular coordinate system of the tablet PC screen proper. Applying this transformation to the coordinates of a token detected in the Webcam image gives the coordinates at which to draw its digital halo on the tablet PC screen, such that the halo appears beneath and around the detected token as it rests on the interactive surface.

[1] A six step process is used: (1) smooth the brightness of the image to remove illumination gradients (discussed in Section A.6.3); (2) perform edge detection by convolving the image with Sobel templates and calculating the magnitude of the resulting gradient at each point; (3) threshold, skeletonise and prune the resulting edge response until only edges of single pixel width remain; (4) find all connected loops of pixels within the edge data that are over a certain length; (5) look for right angles throughout each loop, using the dot product of vectors formed from suitably spaced pixel samples; and (6) confirm detection of the tablet PC when a pixel loop has four clusters of approximately-right angles, each cluster within some threshold distance of one of the four corners of the Webcam image.

A.6.2 Token Code Generation

The presence or absence of token holes at fixed angular increments inside its circumference define a circular binary code that uniquely identifies the token. The generation of these codes is as follows:

1. Initialise an integer *candidate* to zero, and create an empty set of integer *codes*.
2. If *codes* does not contain any rotation of *candidate*, add *candidate* to *codes*. Increment *candidate* by one and repeat this step until *candidate* exceeds $2^N - 1$, where N is the number of potential token holes.
3. Remove from *codes* all values whose binary representations have M or more zeros in succession. This gives each token a balanced appearance and ensures a good triangulation of the token centre when the token is "decoded" using computer-vision techniques.

In my token collection, $N = 13$ and $M = 4$, giving an address space of 505 token codes. Different regions of this address space are allocated to different token types in such a way that the type of a token can be inferred from its hole-based code. Figure A.5 shows the address space allocated to the task tokens of a single user.

A.6.3 Token Detection and Identification

The source images captured by the Webcam are taken at 320×240 pixels and rendered in full colour, as shown in image (a) of Figure A.6. However, reasonable localisation of tokens can be performed at significantly lower resolutions. As expected, the opaque tokens against the bright background gives a sharp contrast, and experimentation led to the initial processing in the final system being performed on images 1/9 of this size, at 107×80 pixels. Photograph (b) of Figure A.6 shows the low resolution greyscale image derived from the *full resolution (FR) colour* image by sampling every third pixel in both dimensions.

Smoothing of Image Brightness

The problem with a simple greyscale conversion of RGB (Red–Green–Blue) colour values is that the reflectivity of the tablet PC screen, combined with the changing response of the Webcam, can cause gradients in brightness across the captured image. This needs to be treated in order for tokens to appear uniformly dark against a uniformly bright background, which substantially eases the token detection process.

Firstly, an image labelled *low resolution (LR) brightness* is derived from the *FR colour* image, based on a down-sampling using the brightness component of the HSB (Hue, Saturation, Brightness) colour space. It is then necessary to 'smooth' the brightness of this image in such a way that the brightness of tokens is not just relatively different to, but absolutely distinguishable from, the brightness of the underlying surface. A square template is convolved with the *LR brightness* image to calculate the mean brightness in the neighbourhood of each pixel. A new image, the *LR smoothed brightness* image, is then created as follows:

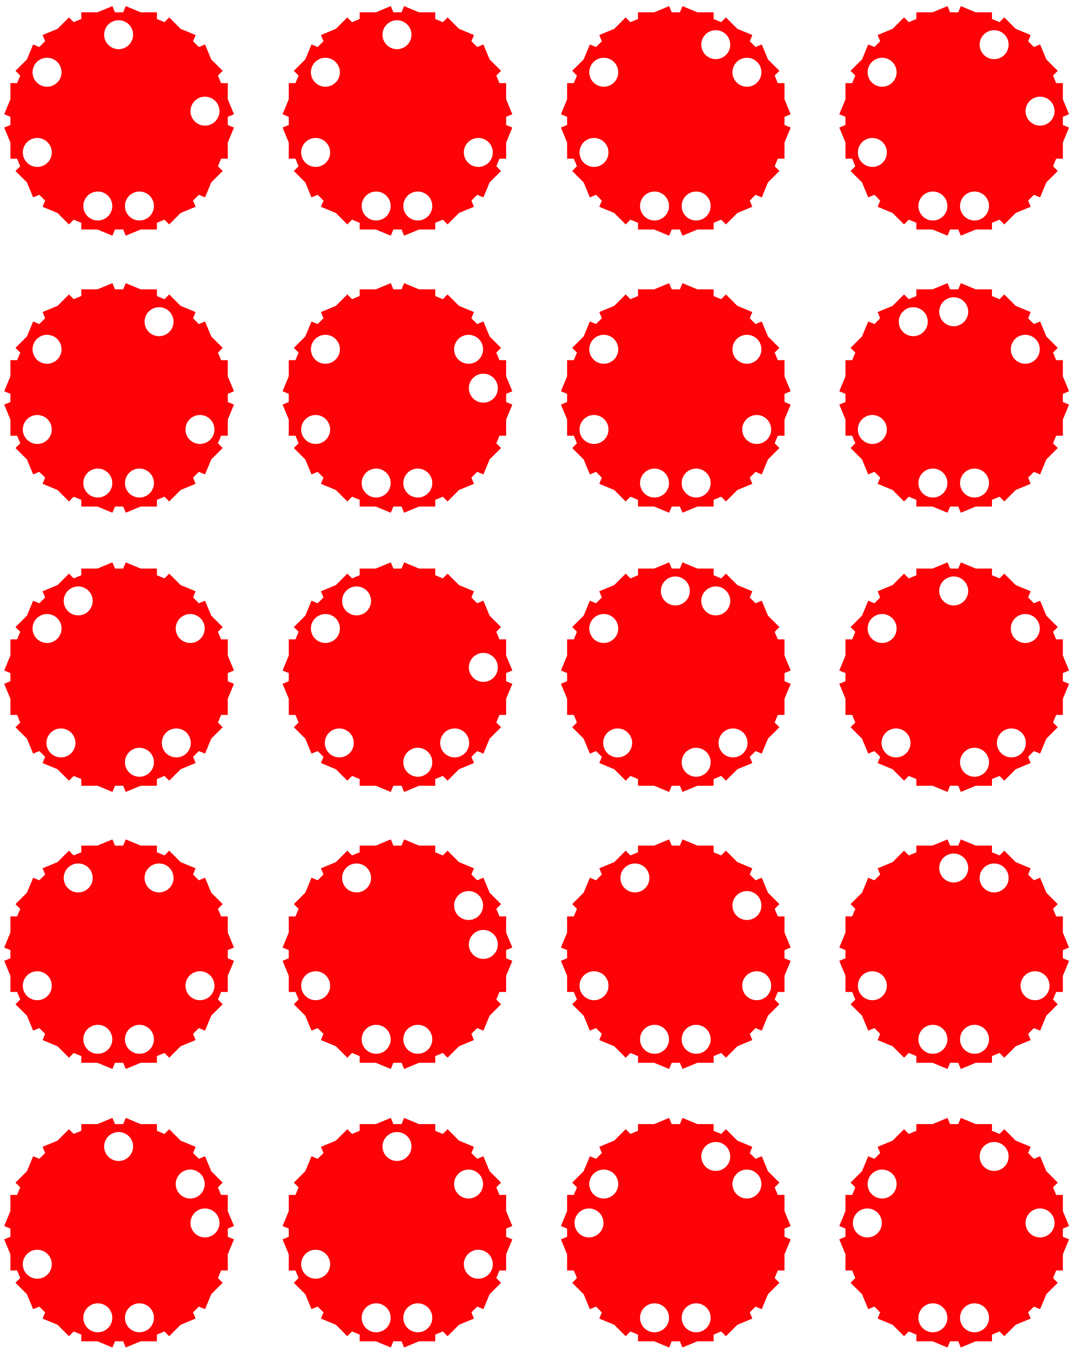

Figure A.5: Sample of Vector Graphics Token Specifications: A set of 20 matching task tokens belonging to a single user, created in Adobe Illustrator based on the output of my "token sheet generator" written in Java. Tokens were laser-cut directly from these specifications.

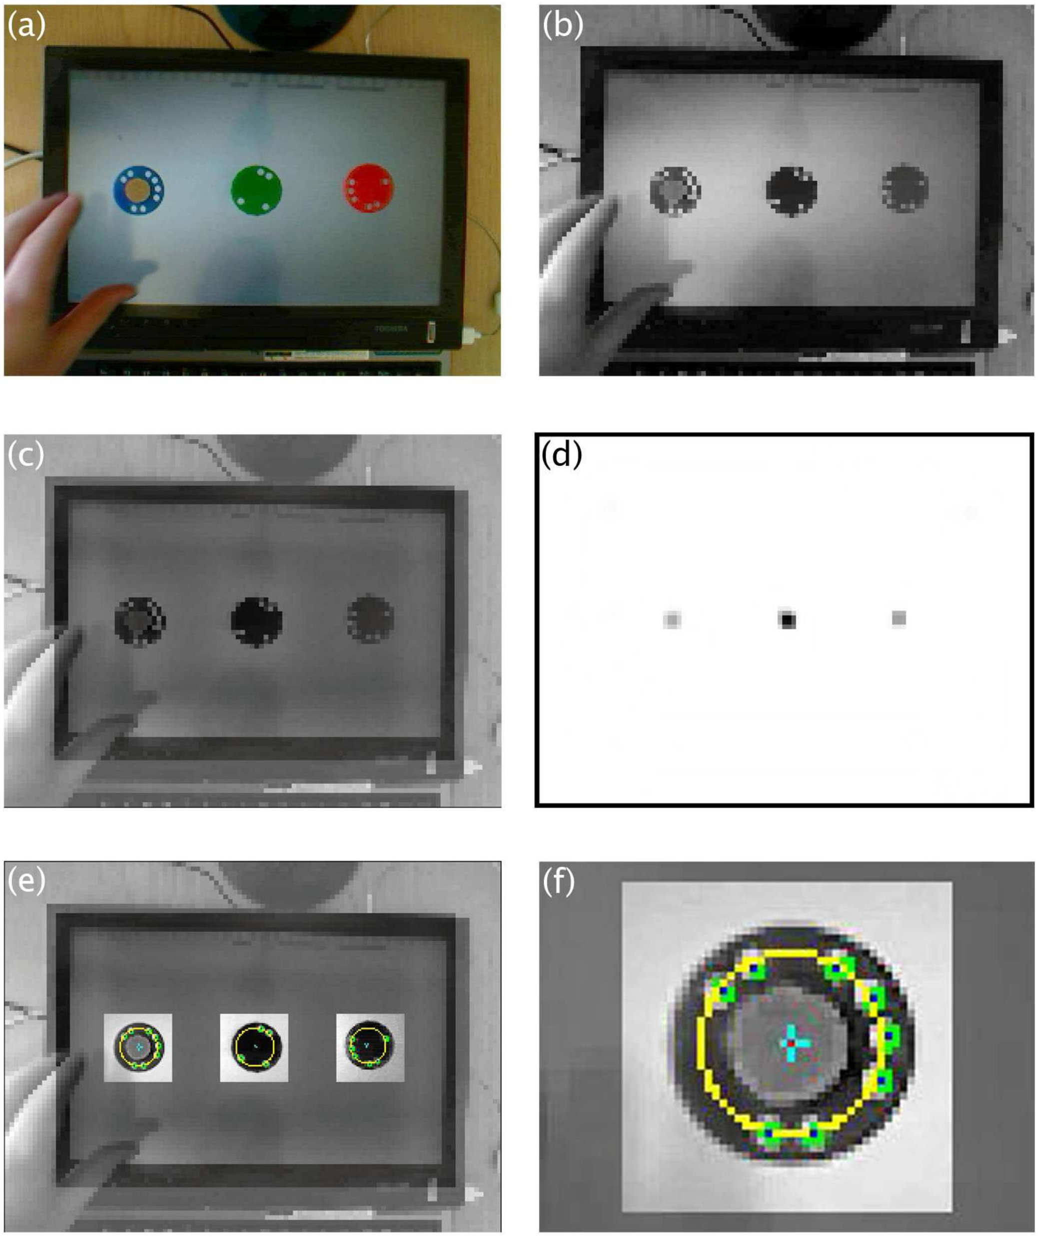

Figure A.6: Token Detection and Identification (1): (a) the high-resolution, colour Webcam image; (b) a low-resolution simple greyscale image (not used); (c) a low-resolution greyscale image based on a weighted combination of saturation and smoothed brightness; (d) response of a fast radial-symmetry "token" detector; (e) high-resolution 'windows' around potential tokens examined for token holes; (f) close-up of token decoding.

$$LR\ smoothed\ brightness\ [x][y]\ = LR\ brightness\ [x][y]\ +\ (0.5\ -\ LR\ means\ [x][y])$$

In regions of higher brightness, where the mean brightness is higher than 0.5, pixels in the local neighbourhood will have there brightness reduced. Conversely for regions of lower brightness, pixels will have their brightness increased. Overall, the effect is to manipulate pixel values such that the mean brightness in any given neighbourhood is 0.5, and this is consistent across the image.

Incorporation of Saturation Information

In addition to brightness information, the saturation of coloured pixels is also utilised through the generation of a *low resolution (LR) saturation* image. Whereas the predominantly colourless interactive surface has low saturation values, the brightly coloured tokens are almost fully saturated. This information is combined in a weighted sum with the smoothed brightness values as follows:

$$LR\ processed\ [x][y]\ = \alpha \times LR\ smoothed\ brightness\ [x][y]\ +\ (1-\alpha) \times (1 - LR\ saturation)$$

A value of $\alpha = 0.6$ proved to give the best results in a variety of lighting conditions. In Figure A.6, image (c) shows the effect of this processing.

Fast Radial Symmetry Detection of Tokens

The *LR processed* image is the source image for token detection, based on the fast radial symmetry detection algorithm of Loy and Zelinsky (2003). This algorithm intuitively works as follows (please refer to the original paper for details of tuning parameters):

1. *Calculate the gradients in the image.* Convolve the image with Sobel templates to produce two arrays: the magnitudes of the gradient at each pixel in both the x- and y-directions.
2. *Find circles of pixels of specified radius whose gradients all point inwards towards a common centre.* For each radius r of circle to detect, consider each pixel *p[x][y]* in turn. Calculate the direction d and magnitude m of the gradient at *p[x][y]* based on the x and y components, and in a similar rectangular array q, increment by m the pixel *q[x′][y′]* located a distance r away from *p[x][y]* in the direction d.
3. *Combine the individual responses for different circle radii into a single output image.* Perform a square blur of template width $2 \times r$ on the individual response for circles of radius r, and for each pixel of each blurred response *b[x][y]*, increment the single output image by[2] *b[x][y]/r*.

[2] A square blur on a rectangular array z creates another rectangular array z' in which the value at *z′[x][y]* is the average of the values in z within the square that is centred on *z[x][y]*. The purpose of performing a square blur in this instance is to disperse the effects of circle centres found at any one radius such that clusters of the 'same' circle centre merge together. The purpose of dividing by r the contribution of responses at radius r is to ensure that each radius of circle receives the same weighting – without this adjustment, detection would be biased towards larger radii, since they have more pixels along their circumference contributing towards the final output value.

For the detection of tokens in the *LR processed* image, a fast radial symmetry detector is used to look for dark circles at radii of 6, 7, and 8 pixels. In Figure A.6, image (d) shows the response of this detector.

Connected Components Analysis

The continuous response from the output of the fast radial symmetry detector – labelled *low resolution (LR) token response* – is used to identify potential token centres through a connected components analysis. The image is first binarised by thresholding the continuous output, such that of the range of output values, the darkest 20% are translated into 1s in the *low resolution (LR) binary image*. The resulting connected regions correspond to potential tokens, and are detected using 4-connected component analysis, in which pixels are considered connected is they are both of value 1 and adjacent horizontally or vertically, but not diagonally. The centres of these components are used as estimators of token centres in the subsequent token identification process.

Detection of Token Holes

The coordinates of each potential token centre are scaled up to the corresponding positions in the *full resolution (FR) colour* image, and a new image is created – *full resolution (FR) processed token* image – of the potential token and its surrounding area from the *FR colour* image. This image is processed as for the *LR processed* image, but without the brightness smoothing since this was found to interfere with the detection of token holes. Another radial symmetry detector is applied, this time responding to potential token 'holes': bright circles of radii 1, 2, 3, and 4 in the *FR processed token* image. The conversion of the continuous response of the fast radial symmetry detector into point values is the same as with the detection of token centres. Holes detected at the very centre of the image or its extreme corners are rejected, since they are likely to be due to reflections off token attachments or holes from adjacent tokens respectively. Images (e) and (f) of Figure A.6 show this process in action. The green 'blobs' are the connected components where the system has detected token holes, and the dark blue dots represent the centres of these blobs.

Fitting a Circle to Token Holes

Any three non-linear points uniquely define the centre and radius of a "circumcircle" passing through each of the three points. To find out if the set of detected token holes lie on the radius of a circle (as do the holes cut around the circumference of each token), circle parameters are sequentially generated from each subset of three points, for as long as the resulting radii and centre points all fall within certain bounds of one another. If they do not, the potential 'token' is rejected. Otherwise, the mean centre point is calculated for the next stage of the identification process. In images (e) and (f) of Figure A.6, the red dot represents this mean centre point, derived from the light blue samples. The yellow circle is centred on this mean centre point, and is of the mean sample radius.

The token identification process up to this point is repeated for a different token configuration in Figure A.7.

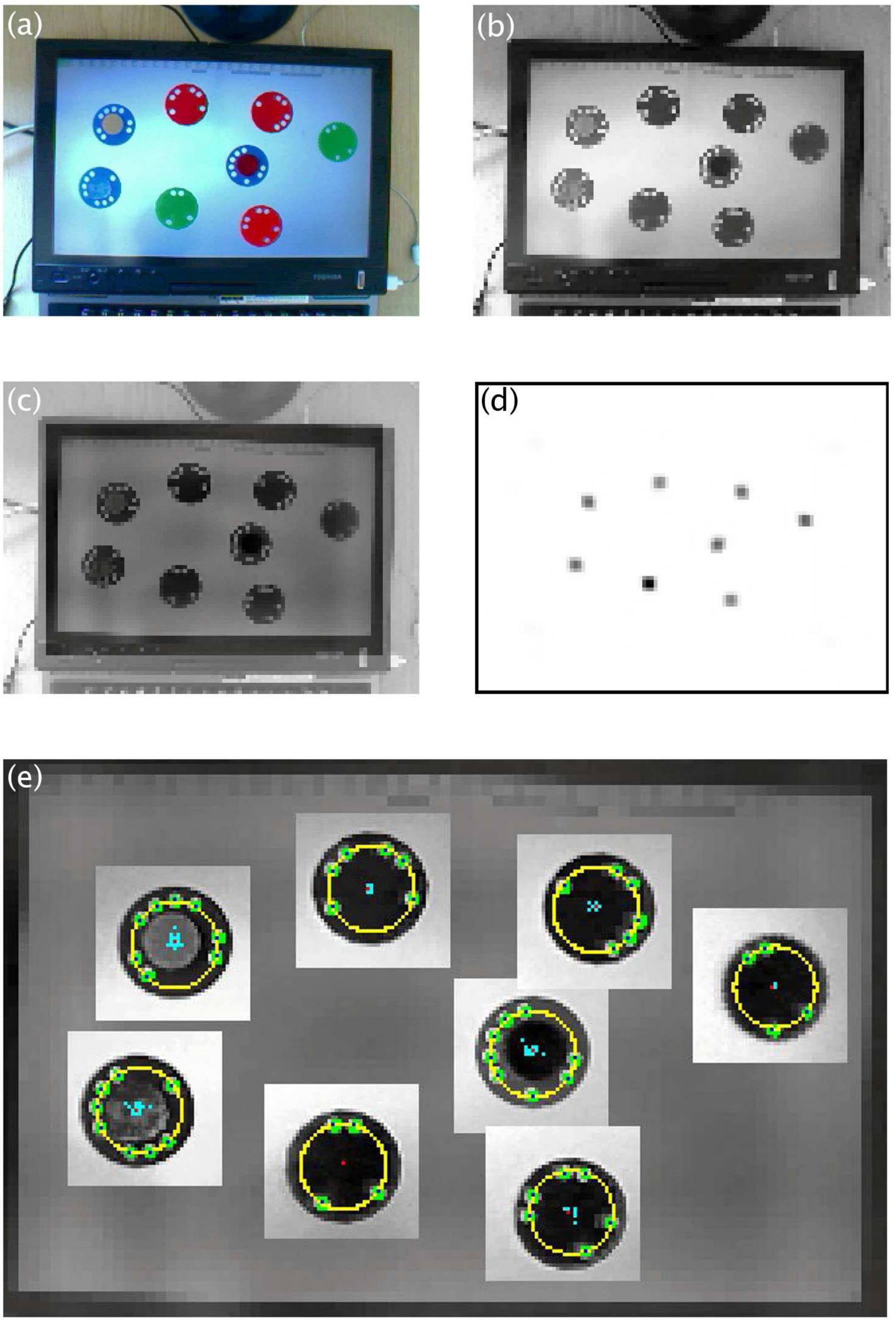

Figure A.7: Token Detection and Identification (2): (a) the high-resolution, colour Webcam image; (b) a low-resolution simple greyscale image (not used); (c) a low-resolution greyscale image based on a weighted combination of saturation and smoothed brightness; (d) response of a fast radial-symmetry 'token' detector; (e) high-resolution 'windows' around potential tokens examined for token holes. Note the harsher brightness gradient across the tablet screen compared with Figure A.6.

Reading the Binary Code

The presence or absence of token holes at fixed angular increments inside its circumference define a circular binary code that uniquely identifies the token. In order to read the binary code, the positions of holes are ordered according to their angle from the centre point, and the angular differences between successive holes are used to determine the resulting binary code. Each rotation of this code is checked against the code pool constructed as described in Section A.6.2, until a match is found or the potential 'token' is rejected.

A.6.4 Token Tracking

In order to provide persistence of token halos across short periods of misdetection, each system representation of a token maintains an internal state machine that is updated on each display cycle of the interactive surface. This requires that on a token's addition on the surface, it must be decoded as the same value for I cycles in succession in order to become "tracked". Tokens are tracked from one cycle to the next by matching the positions of tokens of known identity in cycle c to the positions of the closest token centres in cycle c'. This matching process can also result in new tokens being detected, and tracked tokens being lost (due to being obscured by the hands or taken off the surface). Just as tokens are first *decoded* for I cycles before being tracked, a token can not be *detected* for O cycles in succession before it becomes untracked and its halo disappears from the interactive surface.

Figure A.8 illustrates the progression of token identification and tracking in the fraction of a second following the uncovering of the Webcam lens.

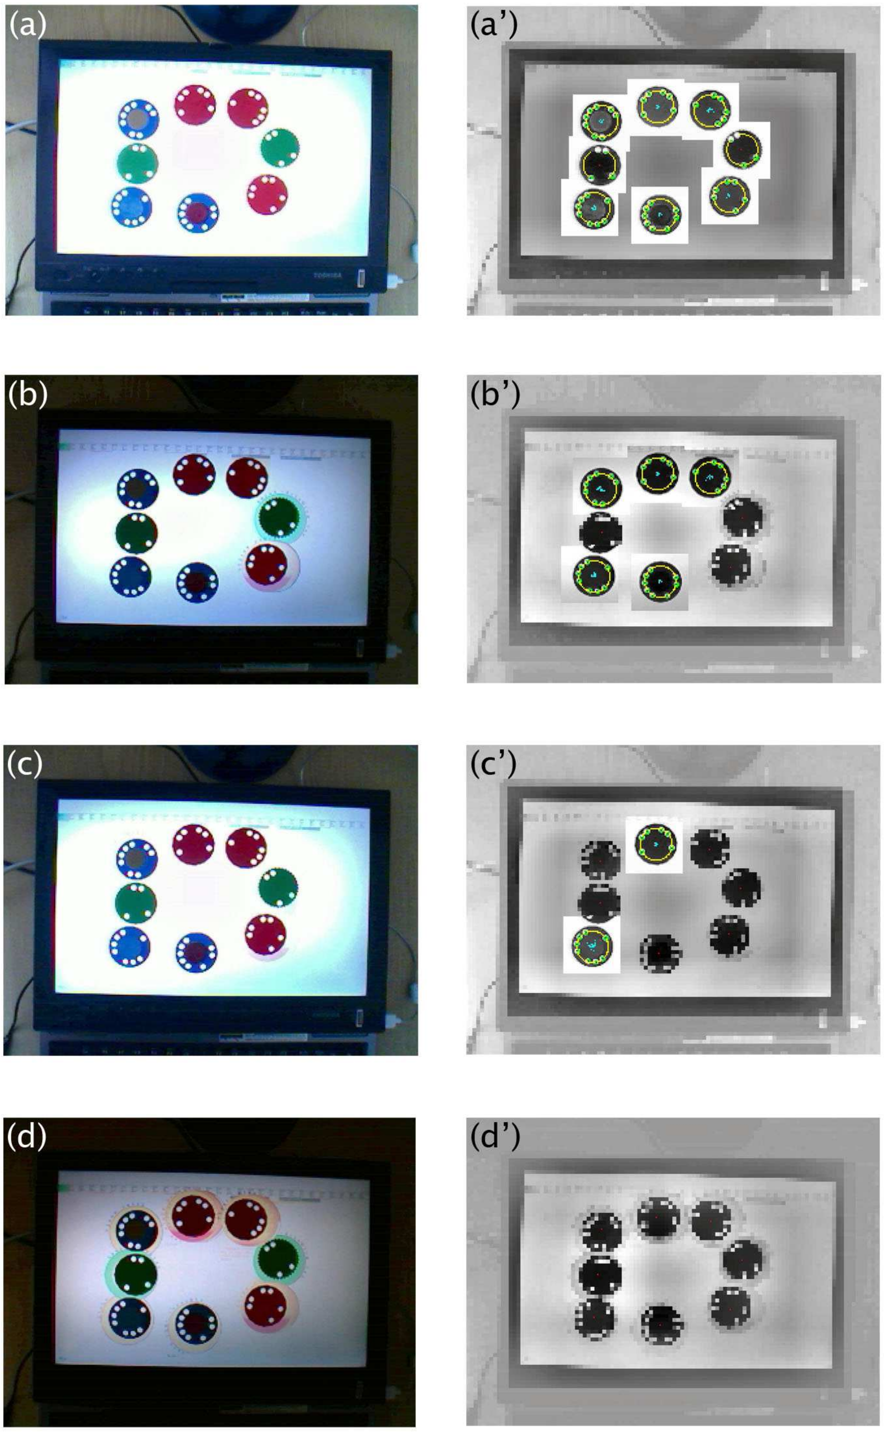

Figure A.8: Token Tracking Using State Machines and Position Matching: (a-d) as the cover is removed from the Webcam, it automatically recalibrates; (a′) all 8 untracked tokens are decoded; (b′) 3 tokens have repeatedly given the same code so can be tracked frame-to-frame based on their position; (c′) 2 token remaining untracked; (d′) all tokens tracked based on their low-resolution position.

A.7 Complete System Guide

COMPLETE GUIDE

INTRODUCTION

This *tangible user interface* uses physical tokens to represent information of common interest to members of small, co-located teams performing project-like work. The purpose of the system is to make routine activities more *efficient* and *engaging*, by allowing users to exploit their *physical environment* and by encouraging *face- to-face conversation*.

THE INTERFACE

Each team member has their own interface consisting of an *interactive surface* and a *control knob*. When tokens are placed on an interactive surface, the surface displays *halos* around each token showing the *attributes* of the information represented. These attributes are selected by *nudging* tokens in the appropriate directions, with the control knob *manipulating* the currently selected attribute. Where *text input* is required, the surface contacts a *server,* which contacts the user's *existing computer*, allowing them to use their keyboard.

TASK TOKENS

These are red tokens, with distinctive *edge textures* used to represent the *owner* of a task. Each user has *20* task tokens. The planned *completion date* and estimated *time remaining* of a task are used to automatically calculate and display its *latest restart time*. The idea is to show when the user would have to work on the task if they left it until the last possible moment, as a means of *visualising and managing workload.* Task tokens also provide a *timer* for task activities, which simultaneously counts down from the estimated time remaining, accumulates *time spent* on the task, and calculates the *percentage complete*. The *actions to do* for a task can also be managed through its token, and tasks can be *delegated* by passing task tokens between team members. The history of a user's tasks can be seen at: http://*server-address*/server/Surface?user=*user-name*

DOCUMENT TOKENS

These are blue tokens shared by all team members, used to create *web-based collaborative editors* for simple text documents. Once created, document tokens can be *cloned* so that access to the document can be shared with other users. In order to access a document, a user *must* have the appropriate document token present on their interactive surface. Once editing, users can specify a unique colour to identify their text, as well as talk to other users through an embedded chat window. To assist team members in finding opportune moments for *ad hoc collaboration*, document tokens show a list of all users currently with the same document token on their interactive surface.

CONTACT TOKENS

These are green tokens representing all users or *contacts* equipped with their own interface. The edge texture of each contact token matches the edge textures of that user's task tokens. When placed on the interactive surface, contact tokens show the current *status* and *work schedule* of that contact. A user's own contact token can be used to set their status and to *recreate* document tokens for previously created documents.

CALENDAR TOKEN

Each user has a smooth green calendar token to *navigate* and *scale* their timeline, and to adjust their global (and individual) *working hours per day* to visualise how this would impact their workload.

PHYSICAL MATERIALS & PROPS

The collection of interfaces is accompanied by a selection of *physical objects and materials* intended for attachment to document tokens as a means of *identification*. Each user also has a special pen with which to *annotate* their tokens. In addition, various *construction materials* and *physical props* are included to provide users with means of *managing* tokens in their physical desktop environment.

Token actions key		Control actions key			
Nudge ↑↓←→	*Nudge in specific direction*	↺/↻	*Turn left/right*	⊙	*Short press*
Nudge ✳	*Nudge in any direction*	⦿+↺/↻	*Press & turn*	⦿	*Long press (~1s)*
Move ●⊂	*Move together*	🖧	*Input from user's computer required*		

TASK TOKENS

Token actions		***Control actions***		
Nudge ✳	*Select task token*	-	-	⦿ *Edit task token* (🖧)
Nudge ←	*Minimise*	-	-	⊙ *Toggle timing*
Nudge ↑	*Select completion date*	↺ *Earlier*	↻ *Later*	⊙ *Toggle timing*
Nudge →	*Select time remaining*	↺ *Less*	↻ *More*	⊙ *Toggle timing*
Nudge ↓	*Select actions to do*	↺ *Down*	↻ *Up*	⊙ *Edit selected* (🖧)
		⦿+↺/↻ *Move down/up*		-
Pass	*Delegate task*	-	-	-

DOCUMENT TOKENS

Token actions		***Control actions***		
Nudge ✳	*Select document token*	-	-	⦿ *Edit document token* (🖧)
Nudge ✳	*Select document token*	-	-	⊙ *Create/Access* (🖧)
Move ●⊂	*Set up document clone*	-	-	⊙ *Clone document token*
Pass	*Share document*	-	-	-

CONTACT TOKENS

Token actions		***Control actions***		
Nudge ↑	*Select status*	-	-	⊙ *Edit status* (🖧)
Nudge ↓ & Move ●⊂$_{Doc}$	*Select created docs*	↺ *Down*	↻ *Up*	-
	Set up document create	-	-	⊙ *Recreate document token*

CALENDAR TOKEN

Token actions		***Control actions***		
Nudge ✳	*Select calendar token*	-	-	⦿ *Jump to today*
Nudge ↑	*Select day pointer*	↺ *Earlier*	↻ *Later*	⊙ *Toggle specific hours*
		⦿+↺/↻ *Hours less/more*		-
Nudge →	*Select normal hours*	↺ *Less*	↻ *More*	-
Nudge ↓	*Select timeline*	↺ *Backward*	↻ *Forward*	-
		⦿+↺/↻ *Zoom out/in*		-

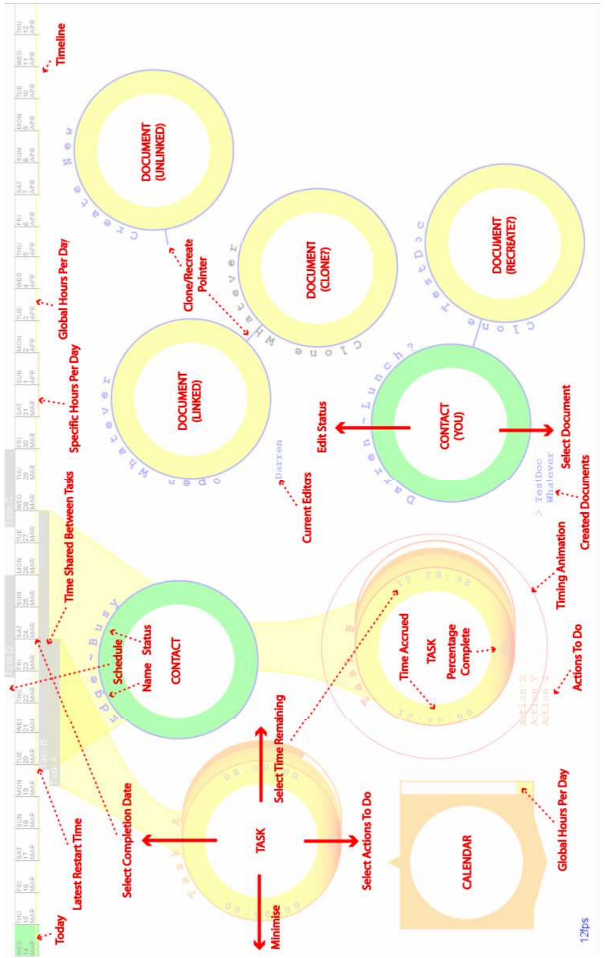

A.8 Gamma Coefficient

Goodman and Kruskal's *gamma* coefficient is a symmetric measure that uses weak monotonicity as its definition of perfect association, and as such is more suitable for the comparison of ranks required here than the more commonly used Spearman's *rho* (Rank Correlation Coefficient), which does not explicitly account for tied data. The *gamma* coefficient is defined as the surplus of concordant pairs (P) over discordant pairs (Q), as a percentage of all pairs ignoring ties[3]), where concordant pairs are those that exhibit positive weak monotonicity (i.e. as x increases, y either increases or stays the same), and discordant pairs are those that exhibit negative weak monotonicity (i.e. as x increases, y either decreases or stays the same). Gamma ranges from -1 to +1, and is interpreted as the proportionate reduction in error from using the relative ordering of two randomly selected untied values in one data set to predict the relative ordering of their respective paired values – the sign of the result indicting the direction of the monotonicity. Table A.2 illustrates the effect of different data pairs on the Gamma Coefficient (X indicates a lack of predictive power due to no untied pairs).

*X*s	*Y*s	*P*	*Q*	*Gamma*
1, 2, 3, 4	1, 2, 3, 4	6	0	1
1, 2, 3, 4	1.5, 1.5, 3.5, 3.5	4	0	1
1, 2, 3, 4	2, 2, 2, 4	3	0	1
1, 2, 3, 4	2, 1, 3, 4	5	1	0.67
1, 2, 3, 4	1.5, 3.5, 1.5, 3.5	3	1	0.50
1, 2, 3, 4	3, 1, 2, 4	4	2	0.33
1, 2, 3, 4	1.5, 4, 1.5, 3	3	2	0.20
1, 2, 3, 4	1, 4, 3, 2	3	3	0
1, 2, 3, 4	4, 3, 2, 1	0	6	-1
1, 2, 3, 4	2.5, 2.5, 2.5, 2.5	0	0	X

Table A.2: Effects of different ranked pairs on Gamma

[3]http://www.statisticssolutions.com/Ordinal-Association.htm

A.9 Derivation of Token Roles

A.9.1 Digital Instrument

To create, inspect, and modify digital information through actions on physical objects.

The application of tokens as instruments used to create, inspect and modify digital information is unique amongst the six usage categories in that it explicitly recognises the computationally augmented nature of the tokens. It is this augmentation that motivates the very existence of tokens in the workplace – the primary design objective of this work – and as such, all of the subsequent categories of use build upon the fundamental premise that the tokens operate within some tangible user interface.

The name DIGITAL INSTRUMENT reflects the primary uses of tokens as both digital representation and physical control, and refers to the ability to make changes to digital information based on visual feedback. In the same way that pilots glance at their instrument panel to check and update their mental models of their aeroplane's status, so too can users of peripheral TUIs take quick glances at the tokens on their surface to update their mental model of the work situation. Similarly, just as pilots need to be able to make fast and frequent changes in response to changing flight conditions, users of peripheral TUIs also need to be able to make such changes in response to changing work conditions.

In the HomeWatch probe, P2 confirmed this need by describing the interface as having "great potential to let you know where you are relative to your targets and so on at a glance, and to make quick updates". Such quick updates are one of the major advantages of tangible interaction – interaction barriers can be lowered to the extent that what was previously an activity performed in batches once a day or once a week, could afterwards be performed in parallel with other, more important activities. It is likely that the low level of this barrier makes the parallelisation of activities highly sensitive to the interaction mechanics involved and the timings of the interactions. Consider the following extract from P1's interview:

> I have the token for my current task and I can very quickly turn that on and off – and that's really quick – but fiddling with the action list isn't so quick . . . Maybe it was that the virtual information on the screen was just less visible than the physical tokens, because of the reflectivity of the screen . . . Maybe it would be better if there was something more physical, because even if the screen was perfect, at a glance I still wouldn't see it as easily as something physical. I guess I didn't use the actions because they were just less tangible!

Whilst P1 was happy to toggle the timers of tasks when transitioning from one task to another using a nudge-click interaction pattern, she was not happy to manage action lists when transitioning from one action to another using a nudge-click-type interaction pattern. This was also the case with P2 and P3. This could be due to the fact that tokens are seen as a whole using our perceptual ability of gestalt, and that screen-based text is comparatively harder to process preconceptually. Alternatively, it could be due to the added interaction complexity associated with users needing to switch back to their PC for text input, as is the case with action lists. Finally, it could be that there is too much cognitive overhead and chance of distraction to think about names of actions, and to plan in general, whilst in the middle of a task.

A.9.2 Knowledge Handle

To offload information and uncover action possibilities through manual manipulation.

The term KNOWLEDGE HANDLE covers all uses of tokens that involve handling for the purpose of concentration or exploration, using tokens as proxies for the information they represent without making changes to their digital content, in contrast with tokens used as DIGITAL INSTRUMENTS. This distinction is closely related to the distinction between pragmatic and epistemic actions (Kirsh and Maglio, 1994): whereas digital updates make progress towards the user's goal, and can therefore be classified as pragmatic actions, offline handling of tokens in an exploratory way does not make progress towards the user's goal, and can therefore be classified as epistemic. Uses of individual tokens include passing or throwing them from hand to hand, rolling them or turning them over between the fingers, dropping them onto surfaces, sliding and flicking them across the surfaces, spinning them in various orientations, and so on. Users may also play with groups of tokens: spreading stacks of tokens out into lines and then stacking them up again; or moving tokens around between tokens clusters or containers. Users may perform these actions whilst thinking about the information represented by the tokens – using the physical objects to help focus attention – or whilst thinking about something else, using the physical activity of token manipulation as a means of moderating concentration to a level optimal for the task in hand. This is in line with Csikszentmihalyi's (1975) notion of *kinesthetic microflow* – "those activities that involve primarily body movements", for example "touching, rubbing, fiddling with objects", imposing "arbitrary patterns that people use to give shape to their experience". Microflow patterns are the "activities that fill gaps in daily routine", and "give structure to experience in the interstices between action patterns dictated by need and social role". According to Csikszentmihalyi, "some coherent, patterned form of experience appears to be necessary to keep the mind from being overwhelmed by randomness".

The tactile properties of tokens are important determinants of users' propensity to 'fiddle' with them. For instance, in the HomeWatch probe, P2 said that he liked playing with his "prickly tokens" – the designated edge texture for his contact token and task tokens – and that the "plain ones weren't as much 'fun"'. The use of tokens as KNOWLEDGE HANDLES can also take place away from the interface in group situations such as meetings: P2 noted that "if you're informal and collaborative, you could take along a pool of blank tokens to play around with", whilst P1 said that for a future set of tasks she and P2 had coming up, "it would be useful to lay them all out and divide them between us – certainly better than just divvying up a list".

A.9.3 Spatial Index

To structure information in the environment using physical props and spatial layout.

The term "index" is used here in the Peircian sense as an indexical sign, that is, something that stands in "dynamical (including spatial) connection both with the individual object, on the one hand, and with the senses of memory of the person for whom it serves as a sign, on the other hand" (Peirce, 1958). As such, tokens can be used as indices throughout the desktop environment, with the locations of tokens relative to other objects acting to draw attention to the actions that led to them being placed in such a way. The tokens thus derive meaning from both the objects that tokens are placed in spatial connection with, and the intentions of the person who put them there. For example, if a user places a document token next to their phone because it needs to be discussed during an upcoming phone conversation, whenever the user looks in the region of the phone their attention will be drawn to the close proximity of the token to the phone, as well as their episodic "senses of memory" associated with the placement act, reminding them of the upcoming phone call and its purpose. Similarly, if a user returns to their desk and finds a task token sitting on their chair, this is an index to the actions of another person who has placed it there for a reason – that the task is important enough for the user to have to remove the task token before sitting down, either because it must be dealt with immediately or because it shouldn't be missed, or both. This kind of behaviour has been frequently observed in office environments in the ethnographic studies of Sellen and Harper (2003): "Often, placing files somewhere visible, such as on a chair or desktop, serves an important role. The mere physical presence of a file serves as a continuous yet relatively unobtrusive reminder of actions to be taken or issues to be attended to for current projects".

This use of tokens as SPATIAL INDICES is not restricted to existing items on the desktop: desktop space itself can be the object of reference. In an arrangement of tokens, each acts as an index to its position in space, which is then interpreted relative to other tokens. Collectively, these may refer to some syntactic relation between the tokens – line up, stacking, clustering, etc. – which draws attention towards the conceptual relationship between them, and the reason for arranging them in such a way.

In the HomeWatch probe, P1 started using spatial organisation strategies from an early stage, reporting that "I've got a row just beyond my tablet which is where I have my fairly frequently used tokens that aren't currently on the tablet". She also said that long-term use of the interface would require more sophisticated organisation, and suggested that the use of "bowls", to differentiate between "almost current tasks", "longer-term tasks", and "unused tasks", would make her more organised. This strategy is akin to the concept of hot, warm and cold documents presented in the Myth of the Paperless Office (Sellen and Harper, 2003). She would use them as "higher level to-do lists" that would enable her to say such things as "OK, that bowl has lots in it, perhaps I should go through them". In contrast, P2 was reluctant to give up too much more deskspace, saying that he would prefer some form of vertical organisation, such as a pinboard or rack, that "wouldn't use up too much real estate". Despite both users wanting better ways to organise their tokens using auxiliary physical structures, and a large set of Technic-Lego being available to create such structures, neither user made the time investment of building any organisational apparatus. This could be due to the fact that they weren't quite at the stage of disorganisation that would spur them to reorganise, that the number of tokens in each 'category' was insufficient to merit its own presentation device or container, that they didn't want to be 'judged' based on their creations, or that they didn't have, or want to be seen to have, the time to spare 'playing' with Lego. In any case, future TUIs based on the concept of peripheral interaction should probably include specialised organisational apparatus, or at least prefabricated Lego-like constructions, in addition to boxed construction kits.

In the HomeWatch probe, all users spoke about the benefits of using tokens as reminders. The following comment by P1 is representative of other comments made about the use of tokens to remind and to refresh (emphasis added):

> I have had the tasks that are the things that I need to do in the next few days on my surface, therefore they're sitting there, even if they're not the currently active task, they're *reminding* me of the things I need to think about . . . But again, I look at the surface often enough just to *refresh my mind* about the tokens that are on there.

Tokens can also be used independently of any digital coupling, as highlighted by P2: "tokens don't have to exist as virtual things, but having them as physical things works well as a reminder". He also thought that tokens could be used to deal with interruptions, saying that they "help when you're half way through [a task] and have to come back to it later".

A.9.4 Material Cue

To aid recognition and reduce search costs by exploiting graphical and material variables.

Tokens can not only be configured in space, but also in terms of their graphical and material variables. The term MATERIAL CUE refers to the way in which these variables can be tailored to increase the visual salience and memorability of tokens. Graphical variables are those that could be represented on paper or a screen, such as object size, shape, and colour. Material variables, on the other hand, are things unique to the physical world: texture, weight, density, malleability, and so on. In this prototype system, some of these were fixed – token size, colour, edge texture, and the presence of a recess – whilst others were available for user customisation through annotation of token surfaces, and attachment of physical objects and materials.

In her initial interview, P1 brought up the fact that token adornment has the benefit of improving recognition:

> I think they reduce the time it takes when I glance at the tablet; it takes me whole seconds to actually read the words on the token surfaces, whereas with sticky things I can instantly think, "Ah, the fuzzy one, that is. . . " whatever. That really helps me. I think the combination of different textures, colours and objects really reduces the time you have to take peering at the tokens.

The extent of P1's perceived benefit of token adornment was such that she attached objects and materials to many of her tokens that did not have a recess: the contact tokens of herself and P2, as well as some long-term and ongoing tasks. In the following interview excerpt, P1 described the rationale behind her material mappings:

> My "usability" document has a red fuzzy thing on top because usability is just a bit fuzzy. My calendar has a white thing with the appearance of a grid that looks like the grid of a calendar. [P2] has a cog because he takes things apart. My "Meta" task, which incorporates all kinds of random stuff, has a ball-bearing because that's a "meta-type" symbol. Even for my contact token, I have a piece of wood with "me" written on it, which is easier than reading just the word "me" from the surface of the token, because it makes it instantly spottable.

In terms of token augmentation, P3 said that "I think the material was a good choice too, because you can write on it and wipe off easily, and you can also stick things onto it", whilst P2 noted "how easy" it was to take a token and write something on it.

Token adornment also had advantages in terms of interaction. For instance, P1 described the "sphere" attached to her "Meta" token as acting like "a kind of handle with which I can pick the token up and move it around; it's really quite satisfying!". Similarly, P2 attached a magnetic rod to his calendar token that allowed him to "pick it up easily and place it accurately on a 'busy' surface". Token augmentation also provides an "aid to identification" and an "aide memoir" according to P2, and thought that the ability to give "badges" to his task tokens to identify them as part of a group was useful. Finally, P3 thought that tokens were "a way of things not getting lost" – when challenged about the potential to lose tokens, she replied that "Something physical is always easier to see and to find than something that is just in your head".

A.9.5 Conversation Prop

To provide shared representations that support discussion and make outcomes explicit.

Tokens can be used in meeting situations to both define the broad topic of discussion, and to support focused attention through verbal deictic references – referring to "this" or "that" in concert with gestural or pointing actions. The concreteness of the shared physical representation can also reduce ambiguity relative to discussion alone, whilst the flexibility of token manipulation, rearrangement and reconfiguration can also be brought to bear on meeting situations within the CONVERSATION PROP usage type, easing the communication of ideas between co-located individuals.

Tokens can also be used to make conversational outcomes explicit, as indicated by P2 in the HomeWatch probe: "the outcome of the meeting could be who walks away with what tasks". P3 highlighted another benefit of such "tangible planning" – no single person is reliant on writing up the outcome of the meeting, meaning that "each person has a set of tokens that they've written on, and they've only got themselves to blame if they lose them". She also commented on the benefit of tokens as means of initiating conversation:

> Tokens can give you an excuse to go over to people's desks and speak to them, and it would help you, when you're doing a task that someone has given you, that you've had a few minutes when they've handed the token over to you to confer and get confirmation of what you're both working on. Tokens actually encourage that person to walk over and talk to you.

This use of tangibles for the initiation of conversation is a phenomenon that has already received some interest in the literature. Brewer et al. (2007), in their study of awareness mechanisms for collaborative groups, report that "several group members expressed an interest in objects or displays that could serve as 'talking points' for newcomers, though at the same time they did not wish for distractions from their day-to-day work".

A.9.6 Social Currency

To indicate work roles and relationships through ownership, possession and exchange.

Physical tokens can stand for a number of different things aside from their digital augmentation, including the representation of social roles and relationships. The ownership of a token can indicate power and authority, manifest in the exertion of influence through its display and exchange. Token possession can indicate rights or responsibilities of the holder, and if the token belongs to someone else, it can also represent a relationship of agency or expectation respectively. The use of tokens in this way must be learnt through participation in the social structures that tokens make explicit, using existing institutional knowledge to interpret the new appropriations, in terms of social symbolism, that tokens make possible. As such, tokens can be seen as a SOCIAL CURRENCY – an institutional fact that establishes the value of tokens as a means of accessing information, encouraging social encounters, and externalising social networks and relationship structures.

In the HomeWatch probe, there were multiple references to tokens being used to represent "ongoing responsibilities", as well as their use as a means of "handing-off responsibilities". However, the 'team of two' – P1 and P2 – had very different ideas about the social entailment of token exchange. P1 reported that "[P2] and I enjoy exchanging tokens. We just reach across and hand tokens over – it creates a bit of mystery". In contrast, the following interview excerpt shows P2's view, demonstrating that task delegation by token passing embodies multiple levels of symbolism that can result in much stronger reactions than might be anticipated:

> I dread receiving tokens! Being given a token is a very definite thing; saying, "I've created a token and I'm giving it to you to do", instead of collaboratively saying, "Hmm, do you think we should do this?" – it defines a power relationship. In that respect it's quite hierarchical.
>
> *Do you think that the creation of power relationships is inherent in the passing of the tokens?*
>
> The thing is it's definite, it's a real thing, there's no ignoring it – it's there, it's round, and someone has defined it virtually for you to do.
>
> *Might the circumstances of the exchange affect this?*
>
> Kind of. You can mitigate it to some extent but it's still there. There's no ignoring a bloody token! It's the mark of Cain. You might as well do it and give it back.

The design of any workplace TUI should therefore provide a degree of flexibility to allow people to 'work around' any undesired symbolic implications, with the ability to appropriate a spare set of task tokens as "team tokens" – this would be very important from P2's point of view. He also suggested that tokens could be used as "credits" in a reward-based system, giving the holder legitimate access to leisure resources such as computer games. Any flexibility can be used in a variety of ways, however, including the imposition of further social pressure through 'guilt' symbols as suggested by P1: "What would be really good is if you had a token that represented the current task on the critical path, to get the person responsible for it to focus on getting that task done".

The example of Nimio (Brewer et al., 2007) cited under CONVERSATION PROP is also relevant here, since these wireless devices for symbolic communication are also of distinctive shape and colour, allowing "users to identify themselves with the toy, rather than identifying the toys with a specific office or user". Users are thus given the ability to "constitute the context and negotiate the configuration of that context by means of exchange of the toys", which they can then use to "represent the workplace as a social space rather than a physical one [...] transforming the problem of representing individual activity into one of group flow". Interestingly, the exchange of Nimios did not occur at any point in the study; one user reported "becoming friends" with hers, and in general the unique Nimios came to symbolically represent their 'owners'. However, more akin to the physical exchange of tokens is the use of Nimios as communication devices. In this respect, it made "stronger and weaker ties more visible" in the same way that patterns of token exchange represent underlying social relationships within work groups. An additional quality, which was also reflected on multiple times by all users involved in my study, was that the existence of research prototypes in the office setting was appreciated as a symbolic representation of the cutting-edge, forwards-looking nature of the organisation. A final point of interest is that while the Nimio devices aimed to support awareness within the group, in reality it was not the glowing colours that facilitated this awareness, but the face-to-face interactions that were motivated by the nature of the technology as a topic of small talk, capable of providing an accessible route into deeper conversation. This relationship between SOCIAL CURRENCY and CONVERSATION PROP is therefore a fluid one, changing dynamically and spontaneously with the incidence of social encounters.

Bibliography

Alcosser, E., Phillips, J., and Wolk, A. (1967). *How to build a working digital computer*. 25

Alonso, M. B. and Keyson, D. V. (2005). MusicCube: making digital music tangible. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1176–1179, New York, NY, USA. ACM Press. 39, 47

Bærentsen, K. B. and Trettvik, J. (2002). An activity theory approach to affordance. In *NordiCHI '02: Proceedings of the second Nordic conference on Human-computer interaction*, pages 51–60, New York, NY, USA. ACM Press. 62

Balakrishnan, R. and Hinckley, K. (1999). The role of kinesthetic reference frames in two-handed input performance. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 171–178, New York, NY, USA. ACM Press. 95

Beaudouin-Lafon, M. (2000). Instrumental interaction: an interaction model for designing post-WIMP user interfaces. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 446–453, New York, NY, USA. ACM Press. 59, 60

Bentley, R., Hughes, J. A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D., and Sommerville, I. (1992). Ethnographically-informed systems design for air traffic control. In *CSCW '92: Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, pages 123–129, New York, NY, USA. ACM Press. 75

Bertelsen, O. W. and Bødker, S. (2003). *HCI Models, Theories and Frameworks: Towards a Multidisciplinary Science.*, chapter Activity Theory, pages 103–133. Morgan Kaufman, San Franciso, CA. 182

Bertin, J. (1967). *Semiologie Graphique*. Mouton, Paris/Den Haag. 57

Beyer, H. and Holtzblatt, K. (1998). *Contextual design: defining customer-centered systems*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 32, 51

Blackwell, A., Edge, D., Dubuc, L., Rode, J., Stringer, M., and Toye, E. (2005). Using solid diagrams for tangible interface prototyping. *IEEE Pervasive Computing*, Oct-Dec:35–41. 33, 107

Blackwell, A. and Green, T. (2000). A cognitive dimensions questionnaire. http://www.cl.cam.ac.uk/∼afb21/CognitiveDimensions/CDquestionnaire.pdf (Accessed 31 December 2007). 54

Blackwell, A. and Green, T. (2003). *HCI Models, Theories and Frameworks: Towards a Multidisciplinary Science.*, chapter Notational Systems - the Cognitive Dimensions of Notations framework., pages 103–133. Morgan Kaufman, San Francisco, CA. 50, 53

Blackwell, A. and Hague, R. (2001). AutoHAN: An architecture for programming the home. In *IEEE Symposia on Human-Centric Computing Languages and Environments*, pages 150–157. 43, 45, 61

Blackwell, A. F. (2006). The reification of metaphor as a design tool. *ACM Trans. Comput.-Hum. Interact.*, 13(4):490–530. 32

Blackwell, A. F., Morrison, C., and Edge, D. (2007). A solid diagram metaphor for tangible interaction. In *CHI 2007 Workshop on Tangible User Interfaces in Context and Theory*. 36, 41

Blumer, H. (1969). *Symbolic Interactionism: Perspective and Method*. University of California Press. 21, 68

Boehner, K., Vertesi, J., Sengers, P., and Dourish, P. (2007). How HCI interprets the probes. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1077–1086, New York, NY, USA. ACM Press. 135, 136

Brave, S., Ishii, H., and Dahley, A. (1998). Tangible interfaces for remote collaboration and communication. In *CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work*, pages 169–178, New York, NY, USA. ACM Press. 43, 44, 48

Brewer, J., Williams, A., and Dourish, P. (2007). A handle on what's going on: combining tangible interfaces and ambient displays for collaborative groups. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 3–10, New York, NY, USA. ACM. 44, 46, 225, 226

Buehler, R., Griffin, D., and Ross, M. (2002). *Heuristics and Biases: The psychology of intuitive judgement*, chapter Inside the planning fallacy: The causes and consequences of optimistic time predictions. New York: Cambridge University Press. 165

Buxton, W. and Myers, B. (1986). A study in two-handed input. In *CHI '86: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 321–326, New York, NY, USA. ACM Press. 94

Calvillo-Gámez, E. H., Leland, N., Shaer, O., and Jacob, R. J. K. (2003). The TAC paradigm: unified conceptual framework to represent tangible user interfaces. In *CLIHC '03: Proceedings of the Latin American conference on Human-computer interaction*, pages 9–15, New York, NY, USA. ACM Press. 34, 35

Card, S. K., Mackinlay, J. D., and Robertson, G. G. (1991). A morphological analysis of the design space of input devices. *ACM Transactions on Information Systems*, 9(2):99–122. 21

Card, S. K., Moran, T. P., and Newell, A. (1987). The keystroke-level model for user performance time with interactive systems. pages 192–206. 147

Carvey, A., Gouldstone, J., Vedurumudi, P., Whiton, A., and Ishii, H. (2006). Rubber shark as user interface. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 634–639, New York, NY, USA. ACM. 39, 56

Chang, A., Gouldstone, J., Zigelbaum, J., and Ishii, H. (2007). Simplicity in interaction design. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 135–138, New York, NY, USA. ACM. 33

Chang, A., Resner, B., Koerner, B., Wang, X., and Ishii, H. (2001). Lumitouch: an emotional communication device. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 313–314, New York, NY, USA. ACM. 44, 48

Cobley, P. and Jansz, L. (2004). *Introducing Semiotics*. Icon Books. 64, 65

Cohen, J., Withgott, M., and Piernot, P. (1999). Logjam: a tangible multi-person interface for video logging. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 128–135, New York, NY, USA. ACM Press. 40, 47

Crilly, N., Blackwell, A. F., and J., C. P. (2006). Graphic elicitation: using research diagrams as interview stimuli. *Qualitative Research*, 6(3):341–366. 138, 154

Csikszentmihalyi, M. (1975). *Beyond Boredom and Anxiety*. Jossey-Bass, San Francisco, CA. 176, 222

Desanctis, G. and Poole, S. (1991). Understanding the differences in collaborative systems use through appropriation analysis. In *Twenty-Fourth Annual Hawaii International Conference on System Sciences*, pages 750–757. 136

Dourish, P. (2001). *Where the action is: the foundations of embodied interaction*. MIT Press, Cambridge, MA, USA. 29, 35, 65

Dourish, P. (2006). Implications for design. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 541–550, New York, NY, USA. ACM. 51

Dourish, P. and Bly, S. (1992). Portholes: supporting awareness in a distributed work group. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 541–547, New York, NY, USA. ACM Press. 95

Edge, D. and Blackwell, A. (2006). Correlates of the cognitive dimensions for tangible user interface. *Journal of Visual Languages and Computing (JVLC)*, 17(4):366–394. 33, 41, 54, 57, 196

Engelhardt, Y. (2002). *The Language of Graphics*. PhD thesis, University of Amsterdam. 57, 74

Fishkin, K. P. (2004). A taxonomy for and analysis of tangible interfaces. *Personal Ubiquitous Comput.*, 8(5):347–358. 35, 59

Fitzmaurice, G. W. (1996). *Graspable user interfaces*. PhD thesis, University of Toronto. Adviser-William Buxton. 26, 60

Fitzmaurice, G. W. and Buxton, W. (1997). An empirical evaluation of graspable user interfaces: towards specialized, space-multiplexed input. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 43–50, New York, NY, USA. ACM Press. 27, 45

Fitzmaurice, G. W., Ishii, H., and Buxton, W. A. S. (1995). Bricks: laying the foundations for graspable user interfaces. In *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 442–449, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 26, 197

Foley, J. D., Wallace, V. W., and Chan, P. (1984). The human factors of computer graphics interaction techniques. *IEEE Computer Graphics Applications*, 4(11):13–48. 19

Frei, P., Su, V., Mikhak, B., and Ishii, H. (2000). Curlybot: designing a new class of computational toys. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 129–136, New York, NY, USA. ACM. 42, 48

Gaver, B., Dunne, T., and Pacenti, E. (1999). Design: Cultural probes. *Interactions*, 6(1):21–29. 51, 135

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates. 62

Gorbet, M. G., Orth, M., and Ishii, H. (1998). Triangles: tangible interface for manipulation and exploration of digital information topography. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 49–56, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 40, 72

Green, T. and Blackwell, A. (1998). Cognitive dimensions of information artefacts: a tutorial. http://www.cl.cam.ac.uk/~afb21/CognitiveDimensions/CDtutorial.pdf (Accessed 31 December 2007). 54

Green, T. R. G. (1989). *People and Computers*, chapter Cognitive Dimensions of Notations, pages 443–460. Cambridge University Press. 53

Green, T. R. G. and Petre, M. (1996). Usability analysis of visual programming environments: a cognitive dimensions approach. *Journal of Visual Languages and Computing*, 7:131–174. 53, 54

Grudin, J. (1999). Primary tasks and peripheral awareness: A field study of multiple monitor use. Technical report, Microsoft Research. 95

Guiard, Y. (1987). Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. In *Journal of Motor Behavior*, volume 19, pages 486–517. 86, 94

Harper, R., Hughes, J., and Shapiro, D. (1991). *Studies in Computer Supported Cooperative Work. Theory, Practice and Design*, chapter Harmonious Working and CSCW: Computer Technology and Air Traffic Control. 75

Hinckley, K., Pausch, R., Goble, J. C., and Kassell, N. F. (1994). Passive real-world interface props for neurosurgical visualization. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 452–458, New York, NY, USA. ACM Press. 38, 48

Hinckley, K., Pausch, R., and Proffitt, D. (1997). Attention and visual feedback: the bimanual frame of reference. In *SI3D '97: Proceedings of the 1997 symposium on Interactive 3D graphics*, pages 121–ff., New York, NY, USA. ACM Press. 32, 95

Hodges, S., Izadi, S., Butler, A., Rrustemi, A., and Buxton, B. (2007). ThinSight: versatile multi-touch sensing for thin form-factor displays. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 259–268, New York, NY, USA. ACM. 40

Holmquist, L. E., Redström, J., and Ljungstrand, P. (1999). Token-based acces to digital information. pages 234–245. 28

Horn, M. S. and Jacob, R. J. K. (2007). Designing tangible programming languages for classroom use. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 159–162, New York, NY, USA. ACM. 42, 48

Hornecker, E. and Buur, J. (2006). Getting a grip on tangible interaction: a framework on physical space and social interaction. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 437–446, New York, NY, USA. ACM Press. 30, 31, 60

Hornecker, E., Marshall, P., and Rogers, Y. (2007). From entry to access: how shareability comes about. In *DPPI '07: Proceedings of the 2007 conference on Designing pleasurable products and interfaces*, pages 328–342, New York, NY, USA. ACM. 31, 72

Hurtienne, J. and Israel, J. H. (2007). Image schemas and their metaphorical extensions: intuitive patterns for tangible interaction. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 127–134, New York, NY, USA. ACM. 65

Husserl, E. (1927). Phenomenology. *Encyclopaedia Britannica 1927*. 154

Hutchins, E. (1994). *Cognition in the wild*. MIT Press, Cambridge, MA. 21

Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., Roussel, N., and Eiderbäck, B. (2003). Technology probes: inspiring design for and with families. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–24, New York, NY, USA. ACM Press. 51, 135, 137

Ishii, H. and Ullmer, B. (1997). Tangible bits: towards seamless interfaces between people, bits and atoms. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 234–241, New York, NY, USA. ACM Press. 19, 20, 27, 36, 73, 96, 187

Jacob, R. J. K., Ishii, H., Pangaro, G., and Patten, J. (2002). A tangible interface for organizing information using a grid. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 339–346, New York, NY, USA. ACM Press. 41, 59

Jordà, S., Geiger, G., Alonso, M., and Kaltenbrunner, M. (2007). The reacTable: exploring the synergy between live music performance and tabletop tangible interfaces. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 139–146, New York, NY, USA. ACM. 39, 48, 56, 72

Kabbash, P., Buxton, W., and Sellen, A. (1994). Two-handed input in a compound task. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 417–423, New York, NY, USA. ACM Press. 94

Kaltenbrunner, M. and Bencina, R. (2007). reacTIVision: a computer-vision framework for table-based tangible interaction. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 69–74, New York, NY, USA. ACM. 39, 104

Kirsh, D. (1995). The intelligent use of space. *Artif. Intell.*, 73(1-2):31–68. 67

Kirsh, D. (2001). The context of work. *Human Computer Interaction*, 16:305–322. 21, 67, 181, 184

Kirsh, D. and Maglio, P. P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4):513–549. 67, 222

Klemmer, S. R., Li, J., Lin, J., and Landay, J. A. (2004). Papier-Mache: toolkit support for tangible input. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 399–406, New York, NY, USA. ACM Press. 33

Klemmer, S. R., Newman, M. W., Farrell, R., Bilezikjian, M., and Landay, J. A. (2001). The designers' outpost: a tangible interface for collaborative web site. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 1–10, New York, NY, USA. ACM. 41, 47

Kobayashi, K., Hirano, M., Narita, A., and Ishii, H. (2003). A tangible interface for IP network simulation. In *CHI '03 extended abstracts on Human factors in computing systems*, pages 800–801, New York, NY, USA. ACM. 40, 59

Koleva, B., Benford, S., Ng, K. H., and Rodden, T. (2003). A framework for tangible user interfaces. In *5th International Symposium, Mobile HCI Conference.* 60

Labrune, J.-B. and Mackay, W. (2006). Telebeads: social network mnemonics for teenagers. In *IDC '06: Proceeding of the 2006 conference on Interaction design and children*, pages 57–64, New York, NY, USA. ACM Press. 45

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press, Chicago, IL. 65, 92, 172, 199

Larkin, J. H. and Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11:65–99. 33

LeClerc, V., Parkes, A., and Ishii, H. (2007). Senspectra: a computationally augmented physical modeling toolkit for sensing and visualization of structural strain. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 801–804, New York, NY, USA. ACM. 38, 55, 64

López de Ipiña, D., Mendonça, P., and Hopper, A. (2002). TRIP: A low-cost vision-based location system for ubiquitous computing. *Personal Ubiquitous Comput.*, 6(3):206–219. 39, 104

Loy, G. and Zelinsky, A. (2003). Fast radial symmetry for detecting points of interest. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(8):959–973. 104, 213

MacEachren, A. M. (1995). *How Maps Work: Representation, Visualisation and Design*. Guildford Press. 57

MacIntyre, B., Mynatt, E. D., Voida, S., Hansen, K. M., Tullio, J., and Corso, G. M. (2001). Support for multitasking and background awareness using interactive peripheral displays. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 41–50, New York, NY, USA. ACM Press. 96

MacKay, W. E. (1999). Is paper safer? The role of paper flight strips in air traffic control. *ACM Trans. Comput.-Hum. Interact.*, 6(4):311–340. 75, 80

Malone, T. W. (1983). How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.*, 1(1):99–112. 21

Marshall, P. (2007). Do tangible interfaces enhance learning? In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 163–170, New York, NY, USA. ACM Press. 153

Mazalek, A., Davenport, G., and Ishii, H. (2002). Tangible viewpoints: a physical approach to multimedia stories. In *MULTIMEDIA '02: Proceedings of the tenth ACM international conference on Multimedia*, pages 153–160, New York, NY, USA. ACM. 40

McGee, D. R. and Cohen, P. R. (2001). Creating tangible interfaces by augmenting physical objects with multimodal language. In *IUI '01: Proceedings of the 6th international conference on Intelligent user interfaces*, pages 113–119, New York, NY, USA. ACM Press. 36, 48

McKelvin, M., Nestande, R., Valdez, L., Yee, K.-P., Back, M., and Harrison, S. (2003). SeismoSpin: a physical instrument for digital data. In *CHI '03: CHI '03 extended abstracts on Human factors in computing systems*, pages 832–833, New York, NY, USA. ACM Press. 198

Merrill, D., Kalanithi, J., and Maes, P. (2007). Siftables: towards sensor network user interfaces. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 75–78, New York, NY, USA. ACM. 38

Moher, T., Watson, B., Kim, J., Hindo, C., Gomez, L., Fransen, S., and McEneany, T. (2005). StoryGrid: a tangible interface for student expression. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1669–1672, New York, NY, USA. ACM Press. 40, 48

Nelson, L., Ichimura, S., Pedersen, E. R., and Adams, L. (1999). Palette: a paper interface for giving presentations. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 354–361, New York, NY, USA. ACM. 40, 47

Newton-Dunn, H., Nakano, H., and Gibson, J. (2003). Block jam: a tangible interface for interactive music. In *NIME '03: Proceedings of the 2003 conference on New interfaces for musical expression*, pages 170–177, Singapore, Singapore. National University of Singapore. 39, 48

Norman, D. A. (1988). *The Psychology of Everyday Things.* Basic Books. 27, 62, 63, 67, 73

Norman, D. A. (1991). *Designing interaction: psychology at the human computer interface*, chapter Cognitive Artefacts, pages 17–38. Cambridge University Press, Cambridge, UK. 67

Pangaro, G., Maynes-Aminzade, D., and Ishii, H. (2002). The actuated workbench: computer-controlled actuation in tabletop tangible interfaces. In *UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 181–190, New York, NY, USA. ACM Press. 44, 73

Parnas, D. L. and Clements, P. C. (1985). A rational design process: how and why to fake it. In *Proceedings of the International Joint Conference on Theory and Practice of Software Development (TAPSOFT) on Formal Methods and Software, Vol.2: Colloquium on Software Engineering (CSE)*, pages 80–100, New York, NY, USA. Springer-Verlag New York, Inc. 50

Patten, J. and Ishii, H. (2000). A comparison of spatial organization strategies in graphical and tangible user interfaces. In *DARE '00: Proceedings of DARE 2000 on Designing augmented reality environments*, pages 41–50, New York, NY, USA. ACM Press. 46

Patten, J. and Ishii, H. (2007). Mechanical constraints as computational constraints in tabletop tangible interfaces. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 809–818, New York, NY, USA. ACM Press. 44, 46, 57

Patten, J., Ishii, H., Hines, J., and Pangaro, G. (2001). Sensetable: a wireless object tracking platform for tangible user interfaces. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 253–260, New York, NY, USA. ACM Press. 40, 45, 59

Patten, J., Recht, B., and Ishii, H. (2002). Audiopad: a tag-based interface for musical performance. In *NIME '02: Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1–6, Singapore, Singapore. National University of Singapore. 39, 48

Peirce, C. (1931–1958). *Collected Papers of Charles Sanders Peirce, 8 vols.* Harvard University Press, Cambridge, MA. 64, 222

Perlman, R. (1976). Using computer technology to provide a creative learning environment for preschool children. MIT Logo Memo #24. 42

Piper, B., Ratti, C., and Ishii, H. (2002). Illuminating clay: a 3-d tangible interface for landscape analysis. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 355–362, New York, NY, USA. ACM. 36, 55, 63, 74

Polynor, R. (1995). The hand that rocks the cradle. *I.D.*, pages 60–65. 38, 45

Poupyrev, I., Nashida, T., Maruyama, S., Rekimoto, J., and Yamaji, Y. (2004). Lumen: interactive visual and shape display for calm computing. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Emerging technologies*, page 17, New York, NY, USA. ACM. 44, 55

Pousman, Z. and Stasko, J. (2006). A taxonomy of ambient information systems: four patterns of design. In *AVI '06: Proceedings of the working conference on Advanced visual interfaces*, pages 67–74, New York, NY, USA. ACM. 96

Preece, J., Rogers, Y., and Sharp, H. (2002). *Interaction design: Beyond human-computer interaction*. John Wiley & Sons, New York. 21, 67

Raffle, H., Parkes, A., Ishii, H., and Lifton, J. (2006). Beyond record and play: backpacks: tangible modulators for kinetic behavior. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 681–690, New York, NY, USA. ACM. 42, 47

Raffle, H. S., Parkes, A. J., and Ishii, H. (2004). Topobo: a constructive assembly system with kinetic memory. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 647–654, New York, NY, USA. ACM Press. 41, 47, 56, 59, 61

Rauterberg, M., Fjeld, M., Krueger, H., Bichsel, M., Leonhardt, U., and Meier, M. (1998). Build-it: a planning tool for construction and design. In *CHI '98: CHI 98 conference summary on Human factors in computing systems*, pages 177–178, New York, NY, USA. ACM Press. 41

Reitmayr, G., Eade, E., and Drummond, T. (2005). Localisation and interaction for augmented maps. In *IEEE ISMAR'05*. 36, 55, 59

Rekimoto, J., Ullmer, B., and Oba, H. (2001). DataTiles: a modular platform for mixed physical and graphical interactions. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 269–276, New York, NY, USA. ACM Press. 38, 48

Reuleaux, F. (1876). *The Kinematics of Machinery: Outline of a Theory of Machines*. Reprinted by Dover Publications in 1963. 196

Richter, J., Thomas, B. H., Sugimoto, M., and Inami, M. (2007). Remote active tangible interactions. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 39–42, New York, NY, USA. ACM. 44, 47

Robson, C. (2002). *Real World Research (Second Edition)*. Blackwell. 135, 138

Rogers, Y. (2006). Moving on from Weiser's vision of calm computing: Engaging UbiComp experiences. In *Ubicomp*, pages 404–421. 20, 153

Ryokai, K., Marti, S., and Ishii, H. (2004). I/O brush: drawing with everyday objects as ink. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 303–310, New York, NY, USA. ACM. 38, 47

Sanders, M. S. and McCormick, J. (1987). *Human Factors in Engineering and Design (6th ed)*. McGraw-Hill, New York, NY. 93

Scaife, M. and Rogers, Y. (1996). External cognition: how do graphical representations work? *Int. J. Hum.-Comput. Stud.*, 45(2):185–213. 67

Schön, D. (1987). *Educating the reflective practitioner*. Jossey-Bass Publishers, San Francisco. 33

Sellen, A. J. and Harper, R. H. (2003). *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA. 77, 223

Shaer, O. and Jacob, R. J. K. (2006). A visual language for programming reality-based interaction. In *VLHCC '06: Proceedings of the Visual Languages and Human-Centric Computing (VL/HCC'06)*, pages 244–245, Washington, DC, USA. IEEE Computer Society. 35

Shaer, O., Leland, N., Calvillo-Gamez, E. H., and Jacob, R. J. K. (2004). The tac paradigm: specifying tangible user interfaces. *Personal Ubiquitous Comput.*, 8(5):359–369. 34

Shneiderman, B. (1987). Direct manipulation: A step beyond programming languages. pages 461–467. 19, 61

Sim, J. (1998). Collecting and analysing qualitative data: issues raised by the focus group. *Journal of Advanced Nursing*, 28:345–52. 135

Singer, A., Hindus, D., Stifelman, L., and White, S. (1999). Tangible progress: less is more in somewire audio spaces. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 104–111, New York, NY, USA. ACM. 40

Spiessl, W., Villar, N., Gellersen, H., and Schmidt, A. (2007). VoodooFlash: authoring across physical and digital form. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 97–100, New York, NY, USA. ACM. 39, 47

Stringer, M., Toye, E. F., Rode, J. A., and Blackwell, A. F. (2004). Teaching rhetorical skills with a tangible user interface. In *IDC '04: Proceeding of the 2004 conference on Interaction design and children*, pages 11–18, New York, NY, USA. ACM. 40, 47

Suchman, L. A. (1987). *Plans and situated actions: the problem of human-machine communication*. Cambridge University Press, New York, NY, USA. 21

Suzuki, H. and Kato, H. (1993). AlgoBlock: A tangible programming language, a tool for collaborative learning. In *4th European Logo Conference*, pages 297–303. 42, 59

Taylor, A. S., Harper, R., Swan, L., Izadi, S., Sellen, A., and Perry, M. (2007a). Homes that make us smart. *Personal Ubiquitous Comput.*, 11(5):383–393. 38

Taylor, S., Izadi, S., Ozenc, K., and Harper, R. (2007b). VideoPlay: Playful and social editing of video using tangible objects and multi-touch interaction. In *In adjunct proceedings of IEEE Conference on Horizontal Interactive Human-Computer Systems. Tabletop 2007.* 42

Ullmer, B. and Ishii, H. (1997). The metaDESK: models and prototypes for tangible user interfaces. In *UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology*, pages 223–232, New York, NY, USA. ACM Press. 72

Ullmer, B. and Ishii, H. (2000). Emerging frameworks for tangible user interfaces. *IBM Syst. J.*, 39(3-4):915–931. 28, 63, 90

Ullmer, B. and Ishii, H. (2001). *Human-Computer Interaction in the New Millennium*, chapter Emerging Frameworks for Tangible User Interfaces, pages 579–601. 28, 34, 58, 73, 187

Ullmer, B., Ishii, H., and Glas, D. (1998). mediaBlocks: physical containers, transports, and controls for online media. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 379–386, New York, NY, USA. ACM Press. 38, 55, 72

Ullmer, B., Ishii, H., and Jacob, R. J. K. (2005). Token+constraint systems for tangible interaction with digital information. *ACM Trans. Comput.-Hum. Interact.*, 12(1):81–118. 63

Ullmer, B. A. (2002). *Tangible interfaces for manipulating aggregates of digital information*. PhD thesis. Supervisor-Hiroshi Ishii. 34, 55

Underkoffler, J. and Ishii, H. (1998). Illuminating light: an optical design tool with a luminous-tangible interface. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 542–549, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co. 36

Underkoffler, J. and Ishii, H. (1999). Urp: a luminous-tangible workbench for urban planning and design. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 386–393, New York, NY, USA. ACM Press. 36, 48, 56, 64, 72, 74

Villar, N. and Gellersen, H. (2007). A malleable control structure for softwired user interfaces. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 49–56, New York, NY, USA. ACM. 39

Villar, N., Lindsay, A. T., and Gellersen, H. (2004). Pin & Play & Perform: a rearrangeable interface for musical composition and performance. In *NIME '05: Proceedings of the 2005 conference on New interfaces for musical expression*, pages 188–191, Singapore, Singapore. National University of Singapore. 39

Waldner, M., Hauber, J., Zauner, J., Haller, M., and Billinghurst, M. (2006). Tangible tiles: design and evaluation of a tangible user interface in a collaborative tabletop setup. In *OZCHI '06: Proceedings of the 20th conference of the computer-human interaction special interest group (CHISIG) of Australia on Computer-human interaction: design: activities, artefacts and environments*, pages 151–158, New York, NY, USA. ACM Press. 38, 47

Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3):66–75. 25

Weiser, M. and Brown, J. S. (1995). Designing calm technology. Xerox PARC. 20, 26, 150

Williams, A., Kabisch, E., and Dourish, P. (2005). From interaction to participation: Configuring space through embodied interaction. In *Ubicomp*, pages 287–304. 30

Zhang, J. and Norman, D. A. (1995). A representational analysis of numeration systems. *Cognition*, 57:271–295. 67

Zigelbaum, J., Horn, M. S., Shaer, O., and Jacob, R. J. K. (2007). The tangible video editor: collaborative video editing with active tokens. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 43–46, New York, NY, USA. ACM. 42, 47, 72

Zuckerman, O., Arida, S., and Resnick, M. (2005). Extending tangible interfaces for education: digital Montessori-inspired manipulatives. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 859–868, New York, NY, USA. ACM Press. 41