3 Methodology

As can be seen from the literature review the flow of shapes from two-dimensional to three-dimensional design representations generally, or from virtual to physical three-dimensional design representations specifically, is typically one-way. While sketches may be used to create physical models, or CAD to create rapid prototyped objects, any emergent shapes generated by working on these latter representations are rarely transferred back into the former. It was to address how these emergent shapes might be incorporated into physical two-dimensional, or virtual three-dimensional representations (to enable all classes of three-dimensional representations to be incorporated into form generation more fully), that the initial focus of this enquiry was on examining the *existing* practices employed by designers to transfer shapes between differing design representations (in order to understand the nature of these transformations more fully).

The first data collection strategy, embodied in the reverse-engineering task, took a direct approach to gaining this understanding by examining the period after some form of physical three-dimensional representation has been created (and modified). It investigated the possibility that the apparent acculturation of digital practices in design, suggested by the literature review, might already have produced methods to apply the emergent shapes generated by working with physical three-dimensional representations to their virtual three-dimensional counterparts.

The second data collection strategy, the series of design protocols, also investigated the flow of shapes between representations, and the methods that designers use to transfer shapes between them, but located at the earliest stages in the creation of a design instead. Here the area of interest was in examining longer established methods which have proven effective in transferring shapes between representations, such as when designers move from producing drawings to making three-dimensional representations.

The initial findings from both these tasks, however, uncovered an unexpected (and distinct) phase between the creation of the participants' two-dimensional and three-dimensional representations, and vice-versa. It was the ambiguity of the space that shapes were transformed in this phase that caused a re-evaluation of the aims and purpose of the enquiry.

While the literature review revealed that enabling virtual and physical three-dimensional representations to be used iteratively (in the earliest stages of design) is a worthwhile endeavour with a definite practical aim, it also revealed that the whole area of employing three-dimensional representations in design is little understood and generally under-theorised. It seems to the author that such statements as are made in the literature about the relative efficacy of drawing for shape generation, and making for form development, are less than meaningful before we understand what relationship exists between the shapes generated by drawing and those that are eventually embodied in a three-dimensional design representation. The enquiry therefore came to focus on understanding the initial representational transition from two dimensions to three by an analysis of the design protocols instead. This chapter therefore, as well as describing the reasoning behind, and justification, for the data collection methods (as well as how they were developed) also describes why the enquiry eventually came to focus on an in –depth analysis of the design protocols alone.

The first section examines the basis for knowledge of non-verbal phenomena, such as the spatial reasoning typically employed in form generation, while the second presents an overview of the data collection exercises developed to uncover evidence of that spatial reasoning. The third section discusses how the initial findings from both of these exercises led to the change of emphasis of the enquiry, from investigating information transfer between physical and virtual representations to a fine-grained analysis of the design protocols instead. The fourth section provides an overview of this analysis. It examines how design activity is identified, how it is described and how that description can be used to answer the research questions posed in the opening chapter. The fifth and final section concludes this chapter with a discussion of how the data collected by the protocols is laid out in the subsequent protocol chapters.

3.1 Knowledge of non-verbal phenomena

While many enquiries focus on interviewing designers as their primary data collection method, there are profound epistemological shortcomings attached to employing verbalisations in design research. Daley (1982) observes that: ‘Only a relatively small (and perhaps insignificant) area of that system of knowing and conceiving which makes designing possible may be amenable to verbal description.' The observation that designers, in their highly time-constrained daily practice, willingly choose to expend time and effort making geometrical representations strongly suggests that it would be, at best, oxymoronic to insist on them using words alone to explain to others what they would normally explain to themselves through making these representations. If designers could describe three-dimensional form in a sentence or two, would they still continue to draw, or to make other forms of external representations?

From this it can be seen that there are distinct problems with data collection strategies such as 'think-aloud' protocols in design research. Lloyd *et al* (1995) for example, under the heading of 'the mapping of words to thoughts', sound a warning note about simplistically assuming a direct link between designers' thoughts and the words they may use in protocols that employ concurrent verbalizations with designing. Alexander (1964) highlights a similar potential pitfall when using language to describe the structure of design problems. He comments that using words to describe a design problem can place a false construction on it as: '...every problem has its own structure, and there are many different problems, the words we have available to describe the components of these problems are generated by forces in the language, not by the problems, and are therefore rather limited in number and cannot describe more than a few cases correctly'.

One gambit employed by researchers to avoid the problems of concurrent verbalization can be to ask designers to give a retrospective account of their creative actions. Unfortunately, Dorst and Cross (2001) note that designer's retrospective accounts of design activity are not completely reliable. Or, as Pinker (1998) rather more pithily expresses it: '…creative people are at their most creative when writing their autobiographies.' A further complication that can arise when using verbalisations in design protocols, again highlighted by Lloyd *et al* (1995), is that participants may try to '…second guess what the experimenter 'wants to hear'…' when formulating their responses. Lloyd *et al*, enlarge on this discussion of the potential problems of using concurrent verbalizations in think aloud protocols, when they suggest that there may be a basic incompatibility between the kinds of primarily visual thinking employed in design reasoning and the act of verbalization itself.

Evidence that verbal and visual forms of reasoning may be mutually exclusive can also be found outside of design, in clinical treatments for post traumatic stress disorder (PTSD). Chris Brewin, a professor of clinical psychology at University College London, has developed a strategy for PTSD sufferers which encourages them to produce a detailed narrative of the traumatic event that caused their PTSD. Brewin (2005) states that this narrative works by blocking the involuntary recall of their image-based memory of that event as: '...the verbal memory, that they are elaborating and creating in much more detail, interferes with the image-based memory which they have been constantly retrieving as part of the PTSD'. As the memories triggered as part of PTSD are primarily visual in nature and, as Lloyd *et al* suggest, reasoning about form in design is also primarily visual, Brewin's successful treatment of PTSD provides substantial evidence that these forms of reasoning are in fact mutually exclusive.

If verbalizing interferes with the process of designing, what other methods might be used to make design reasoning explicit? If we assume that it is a wholly internal 'black-boxed' cognitive process, then there seems to be little alternative to also assuming that it will remain a mysterious and fundamentally unknowable one. But, if design reasoning is viewed as a product of the interaction of both internal *and* external representations (rather than a wholly internal one) it is possible to look at this interaction as a way of revealing design reasoning.

In effect the combination of the designer and their external representation is a form of distributed cognitive system. Rogers (2004), for example, has observed that distributed cognitive systems, by breaking down the traditional boundaries between where internal (and therefore invisible) cognition stops and external culture starts, it is possible to use the concepts from traditional cognitive theories of the mind but to apply them more reliably to these larger and, more importantly, directly observable systems.

While reasoning of the kind that employs declarative knowledge may be amenable to verbalization, the spatial and visual reasoning employed in design ideation (the primary function of sketching in the larger design process) may not. However, by treating the interaction of designers with their external design representations as a distributed cognitive system, it is possible to uncover this visual and spatial reasoning (and the situated, embodied and tacit knowledge engaged by designers in support of that reasoning), which would not be discernible by purely verbal data collection methods. Both of the data collection exercises therefore focus on how designers interact with their external representations instead.

3.2 Overview of the data collection exercises

While observing designers *in situ* on live projects would produce data that genuinely reflects design practice, disseminating the results of these observations would be problematic. For all practical purposes it is impossible to separate the process, the intended object of study, from its products. Observing live projects therefore inevitably raises the issue of protecting the intellectual property of the participants.

In addition, the complex phenomena that arise from the interaction of participants with a number of physical and virtual design representations, over an extended period of time, has the potential to generate unwieldy volumes of data.

To obviate both of these problems it has been necessary to set up artificial tasks, and to observe these instead. While they are artificial they have, however, been designed to encompass meaningful chunks of practice which are large enough to include the possibility of employing multiple representations of form, while still small enough not to swamp the enquiry in data.

To generate the kind of data required to answer the initial research questions set in the literature review, a more specific set (listed below) was subsequently developed. As was previously noted, the iterative nature of design breaks down when designers move from designing in a single medium, such as drawing, to employing multiple forms of design representations. These more specific questions therefore focused on the nature of the change of from one representation to another in typical design practice:

- What representations do designers use when their choice is not constrained?
- When, how and why do they change representation?
- How do they move from two-dimensional to three-dimensional solutions?
- How do they move from physical to virtual representations?

The reverse-engineering task examines the period after an initial three-dimensional representation has been produced. The design protocols, conversely, examine the period from the inception of a design to the creation of the first three-dimensional representation:

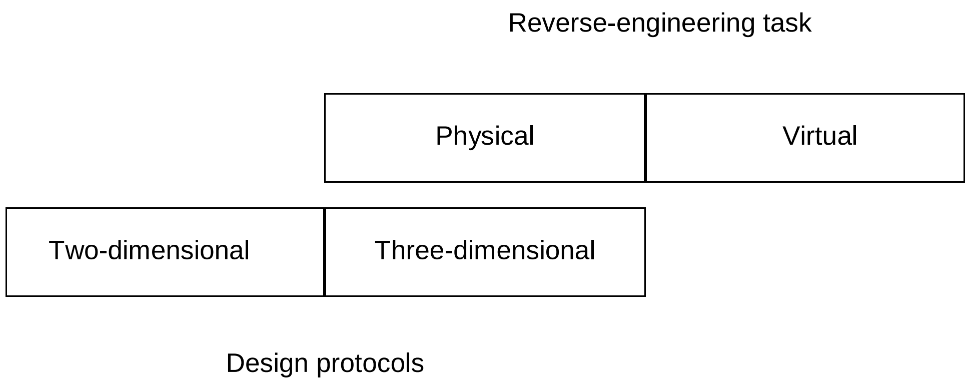

The reverse-engineering task was specifically directed toward addressing the apparent breakdown of the flow of information from physical design representations to virtual ones noted in the literature review.

In contrast to the earlier approaches outlined in the literature review, which employed technologies such as laser-scanning to turn physical geometry into its virtual equivalent, the reverse-engineering task was devised instead to provide a fresh approach, driven instead by contemporary designers' increasing facility with CAD. While technologies such as laser scanning provide an accurate digital representation of the *surface* of a model, that representation is without the kind of *structure* that a designer would put in place if it had been built (from scratch) within a solid-modelling or NURBS surfacing program. Without that underlying structure a virtual model is not easily editable. Although Sutherland (1975) states that the requirement for this structure is part of the objection to CAD modelling as a conceptual tool, it is still significant to the designer. The structure of a CAD representation, like the templates described in the literature review, can be seen as a condensation of the designer's spatial reasoning made explicit in the curves used to generate surfaces.

The second data collection exercise, the series of design protocols, was also directed at examining information flow between representations, but was more general in its scope and purpose than the reverse-engineering task. The design protocols were intended to be more representative of design practice than the reverse-engineering task. To this end they had a formal design brief, specified outcomes and a time constraint, but were completely unconstrained as to the kind of representation that could be employed in them. By allowing participants as much freedom as possible in the choice of representations the design protocols are intended to provide a record of their actions which represents, given the constraints imposed by a PhD enquiry, a sample of their normal form generation practices.

3.2.1 The reverse-engineering task

One of the subsidiary purposes of the reverse-engineering task was to test the thesis that there has been an acculturation of digital practices in contemporary design. The premise underlying it (that there is a burgeoning body of tacit knowledge in contemporary design concerning digital representations) would in turn make it possible for participants to imagine the structures required to build a virtual digital representation after examining a physical three-dimensional design representation.

The reverse-engineering task itself was intended to gather evidence of the kinds of strategies and geometrical information that designers employ to transform their design proposals from physical three-dimensional representations to virtual ones. Being able to transfer information seamlessly and informally between these media would, of course, be a vital component in allowing physical and virtual design representations to be used in an iterative design cycle.

Before arriving at the final version of the reverse-engineering task, however, a number of alternative versions were contemplated. They all shared a common two-part structure which had an assessment element built in to it. In each variant this took the form of a second session where, having previously envisaged the structures required to build a virtual representation from in various media, the participants would then be invited to *use* those structures to build that virtual representation. The only element that varied among these options was the design representation employed.

The first version would have been particularly suited to distance data collection as participants were only required to draw digital structures onto a two-dimensional representation of a three-dimensional object. This was abandoned however due to the interference to that might be caused to the participants' internal three-dimensional visualisations by the two-dimensional nature of the representation employed.

An attempt was made to address this shortcoming in a subsequent variation which specified the use of a three-dimensional physical representation instead. The intended outcome of this particular task was a physical object with the kind of structures required for digital surface modelling marked on it by the participant. There was a possibility, however, that an unintentional bias was also contained in this version of the task as the objects supplied to the participants would have been created on a rapid prototyping machine from a digital model.

A direct and much more intriguing and elegant way of assessing designers' ability to imagine digital structures was also considered in the earlier stages of the enquiry. This required participants to mentally visualise the structures required to make a virtual digital representation while viewing an on-screen representation of a virtual three-dimensional object. Eye-tracking would then have been employed to capture the participants' involuntary eye movements, which are known to directly trace out the same shapes as those they are visualising.

The phenomenon that involuntary eye movements also trace out the contours of areas of interest would, ultimately, have been a very useful tool to visualise the more evanescent and, perhaps unconscious, decompositions that designers apply to the shapes contained in their design representations. As gathering evidence of these internal visualisation processes directly in this way seemed a very fruitful area to the author, it was surprising that, despite making a strong case for this approach in the early stages of the enquiry, it garnered no support. A less direct method of studying some of the more long-lived of these shape decompositions was eventually incorporated into the analysis of the design protocols however. Here they were inferred by examining how designers chose to draw over previously drawn lines, in the subsequent analysis of the sequences of actions observed in the protocols.

A possible shortcoming of employing eye-tracking is that most eye-trackers are distinctly invasive and lab-bound. Even relatively unconstrained systems, such as Tobii's, will only work with on-screen representations. Perhaps there is still a case for working with two-dimensional on-screen representations of three-dimensional forms (that can be rotated on-screen), but it would difficult to ensure that the data collected by this means refers solely to a three-dimensional rather than to a two-dimensional representation.

Although these earlier versions of the reverse-engineering task were intended to produce finished artefacts, such markings on a physical object or a digital file, the decision to video the pilot of the eventual exercise (to make a record that could be examined at leisure to gauge the effectiveness of the methodology employed) revealed so much about the techniques and rationale behind how designers move information between representations that it was decided to record all subsequent ones. The pilot exercise also revealed, against the author’s expectations, that participants might use feature-based modelling as the basis for their digital representations rather than a NURBS-based surface modelling approach. In the light of this it seemed possible to the author that the reverse-engineering task was being over-constrained. Therefore, in the final version of the reverse-engineering task, participants were simply presented with a collection of physical representations of digital designs (shown in Figure 3-1), and were asked to describe how they would create virtual versions of them.

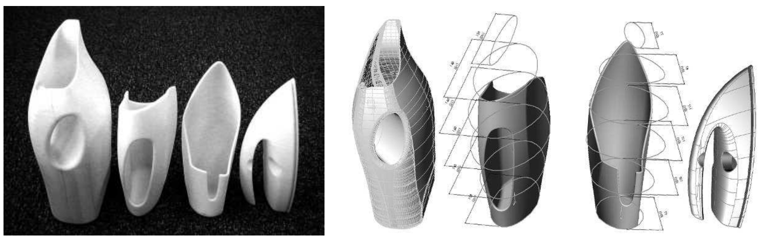

Figure 3-1 the reverse-engineering task objects

The objects themselves were produced through surface modelling (right), feature based parametric solid modelling (centre pair) and a combination of both (left). The centre pair was produced from the same base file, i.e. they had identical digital structures, but were handed in relation to each other as well as having their top profile and the loft profile curves parametrically altered.

3.2.2 The design protocols

The design protocols, in turn, focus on the (relatively) unconstrained process of making representations. There are two main areas of interest in this process: when, how and why do participants switch between representations, and what is the nature of the design ‘moves’ they make within each representation.

Seven individuals took part in a period spanning from Nov 2006 to Mar 2008. They were drawn from an initial pool of eleven individuals which contained both practising designers and design students. No selection of the participants took place other than on the basis of their availability within the timeframe of the enquiry. The seven who actually undertook the protocols were simply those who were available within this timeframe.

The participants were required to generate a design proposal for the handle area of a domestic clothes iron intended to fit onto a pre-existing base element, with an hour being allowed for this task. After signing a consent form they were presented with the design brief, three copies each of three full-size orthogonal views, a perspective view and a physical model of the iron base. Wherever possible the protocols themselves were recorded in the participants' normal working environment. An iron was chosen as the subject of the design brief as it would potentially result in a more organic form, that all participants would be familiar with the subject, and whose design would be equally amenable to both modelling and drawing.

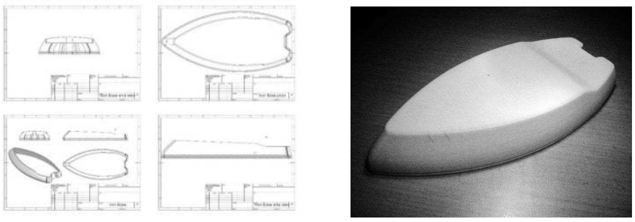

Figure 3-2 the set of drawings and base model supplied to the participants

A significant aspect of the design protocols themselves is that participants were allowed complete freedom in their choice of design representations. This freedom of course raised the possibility that many participants, when faced with one hour in which to generate a design proposal, would only produce drawings in that time. However, rather than presenting the task as one specifically geared to creating physical or three-dimensional representations, participants were instead invited to produce any form of representation, two or three-dimensional, physical or virtual, that would allow some form of *ergonomic* assessment within the stipulated one hour time limit. Although this might be seen as weighting the design protocols toward using physical three-dimensional representations, violating the intention to allow participants complete freedom in their choice of representation, ultimately it proved to be a very successful gambit. Although one participant produced only drawings as a final outcome, while another only produced drawings and a virtual three-dimensional representation, both came up with strategies which allowed them to assess the ergonomic aspects of their designs without creating a physical three-dimensional representation.

The one hour time limit of the protocol may also seem extremely stringent, but this reflects the time constraints which practising designers work under on occasion. It was specifically chosen so that the data generated by the protocol would not be too great to defy analysis within a reasonable space of time, as well as making the task less onerous for participants. The two hours required for each data collection task would of course be very expensive if charged to the researcher at the participants' standard rate. The designers who participate in these protocols on an unpaid basis may be losing valuable time when they agree to take part.

That the one hour allowed for the task in the protocol was representative of the time that might be allowed for the task in normal design practice was borne out by the reaction of the more experienced designers. They saw nothing unusual in being asked to produce a number of design representations in so short a time, and remained quite relaxed throughout the exercise. The reaction of the students, however, was very different, and some were vociferous about being asked to participate in a task that was so time constrained.

3.3 More than two dimensions, less than three

A striking feature that was apparent in all the participants' answers to the design protocols was their choice of predominantly orthogonal views when drawing their initial designs. In addition, when visually assessing the supplied base, participants would often hold it in such a way that it could be viewed orthogonally, sometimes even closing one eye to remove any sense of stereoscopic depth perception. In both data collection exercises, participants also reported using two-dimensional scanning, or taking photographs of orthogonal views of objects that could subsequently be used in a digital modelling program to build virtual geometry over.

Furthermore, in those exercises where participants did go on to make a three-dimensional representation there was also a further ‘orthogonal’ phase in the transition from the initial two-dimensional representations to the final, three-dimensional one. This can be seen in the images below from Andrew’s protocol:

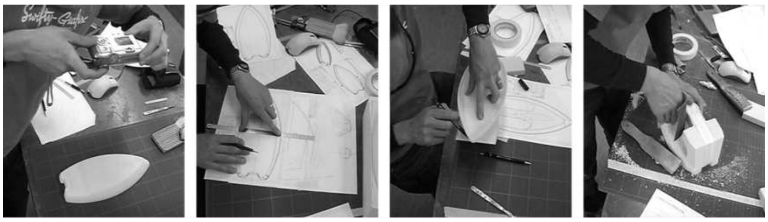

Figure 3-3 using two-dimensional shapes to create three-dimensional objects

A similar tendency to employ orthogonal views was also observed in the reverse-engineering tasks. In the examples in Figure 3-4, taken from Lynn’s task, her first act is to trace around the supplied object to produce orthogonal views of it. In her accompanying commentary she explains that the reason for doing this is because if she were to create a virtual representation of this object it would be created from ‘cut’ features (in a feature-based solid modelling program), which would be derived from digital orthogonal sketches.

Figure 3-4 creating two-dimensional views from three-dimensional objects

What these observations suggest is that the relationship between two-dimensional and three-dimensional design representations is not a direct one. Designers may be employing distinctly orthogonal *two-dimensional* practices when creating *three-dimensional* form. It appears that, in the reverse-engineering task, reducing physical three-dimensional representations to two-dimensional ones enables the participants to extract sufficient design information to be able to re-create them in three *virtual* dimensions. Conversely, in the design protocols, it seems that a similar composite two-dimensional/three-dimensional phase allows designers to transform their two-dimensional drawings into an initial three-dimensional representation when working in the opposite direction.

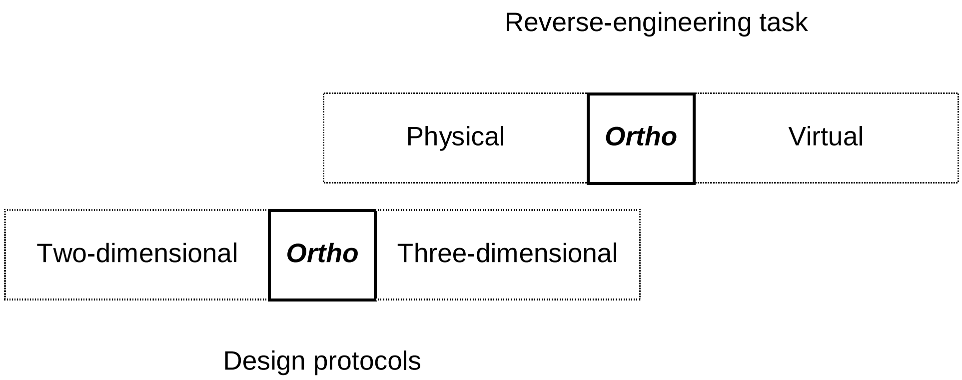

As was stated in the introduction to this chapter, it is the ill-defined spatial nature of the representations employed in this orthogonal phase which led the enquiry away from its original objective, and toward gaining a greater understanding of how this spatially ambiguous orthogonal phase functions in form generation. The method chosen to do this was through a detailed analysis of the four experienced designers' protocols.

3.4 Analysis of the design protocols

While it may be obvious that the physical space occupied by a design drawing is a two-dimensional one, what is not so obvious is the perceptual space inhabited by a designer when producing that drawing. However, as was noted earlier, direct evidence of the perceptual space occupied by designers when drawing (as part of their internal cognitive activity) is problematic to gather. While other forms of reasoning, for example those that employ declarative knowledge, may be amenable to verbalization, spatial reasoning it seems may not. Indeed, Lloyd *et al* (1995) suggest that spatial reasoning may be incompatible with verbalization altogether.

The empirical evidence available to the author before undertaking this enquiry, and based on fifteen years first-hand experience of attempting to convert architectural and industrial designers' sketches into three-dimensional forms, is that the perceptual spaces designers occupy when drawing appear to cover the entire range of spatial reasoning from the two- to the three-dimensional. Although some designers may genuinely be dealing with three-dimensional form in their drawings, others, even though they work in the same medium, and in an apparently identical fashion, may only be engaging with the properties of two-dimensional shape.

Any discussion about the function of drawing in the generation of design proposals therefore may be less than meaningful, in practical terms, before we understand *what* relationship exists between the two-dimensional shapes generated by drawing and the end product of the whole design process. When designers draw, are they drawing two-dimensional shapes, or are they designing three-dimensional forms? Are they responding directly to the requirements of three-dimensional form itself, or are they instead simply responding to the two-dimensional forces within the drawn representation? It was these observations which led to the first of the enquiry's ultimate research questions:

- How three-dimensional *are* the shapes contained in a drawn representation?

As two-dimensional drawings are transformed into three-dimensional representations every day in design practice it might be assumed that the shapes contained in designers' drawings are simply two-dimensional representations of three-dimensional forms. If, as the literature review and the author's experience of design practice suggest, this is not necessarily the case, and some of the shapes contained in designers' drawings do not represent three-dimensional forms, then the need for a further research question emerges. As transforming a drawing into a three-dimensional representation entails mapping from a lower spatial degree to a higher one, where does that additional information come from? The second of the enquiry's ultimate research questions was therefore framed as:

- How are the additional shapes required to create three-dimensional form *generated*?

Schön and Wiggins 'see-move-see' (1992) cycle of design 'moves' is an obvious source of additional information. However, if that additional information is generated in a representation of lower spatial degree from the intended outcome, as happens when designers draw design proposals, will it be the kind required to generate a three-dimensional object? If the shapes generated by drawing are not necessarily three-dimensionally realisable, is it possible, in turn, that the information required to map from two dimensions to three is generated in the process of making *other* design representations instead? Pye's (1978) observation that designing and making are inextricably linked suggests that this is a tenable proposition. However, in order to answer the research questions adequately, it will now be necessary to lay out a complete description of the drawing and making activity contained in each protocol.

Segregating the drawing and making actions contained in each protocol will also provide an opportunity to answer Tovey (2002) and Evan's (2005) criticisms that three-dimensional representations generally are too concrete, depictive and unambiguous to support shape generation effectively. The relative contribution made by drawing and making actions to the generation of a design's three-dimensional form can however be assessed by identifying the type of decompositions employed and the nature of the transformations contained in them. Goel's concept of lateral and vertical transformations is employed to classify which actions qualify as generative ones (lateral transformations), and which qualify as developmental ones (vertical transformations). An additional measure of the flexibility, or 'plasticity', of the various two-dimensional and three-dimensional design representations (particularly as the same representation is often employed in a different fashion by each of the participants) is made by comparing the relative frequency of combinatorial, versus indefinite, decompositions observed during the course of each protocol.

In subsequent protocols, which contain three-dimensional design representations, this distinction between lateral and vertical transformations is also used to identify any possible form *generation* activity in the employment of those three-dimensional representations.

Although it is sometimes assumed, both in design practice and in design theory that the design spaces addressed by drawing are directly concerned with the generation of three-dimensional form, there is no rigorous proof that this is actually the case. While it is obvious that the physical space occupied by a design drawing is a strictly two-dimensional one, what is not readily discernible is the perceptual space that drawings occupy in designers' heads as they undertake drawing or making actions. These drawing and making actions, however, can be mapped against the elemental and transformational space they occur in. This has been achieved by identifying where they reside in Stiny's algebras of shapes ($\boldsymbol{U_{ij}}$), comparing the spatial degree that design elements were transformed in with the spatial degree of the design element itself.

3.4.1 Identifying design activity

Schön and Wiggins design ‘moves’ (1992) have been used as the basis of the fundamental unit of form generation activity employed in the protocols. To avoid loading this fundamental unit with unwanted meanings the deliberately neutral term ‘design action’ is used as a label for it throughout this thesis. Attempting to define a design action by extension, i.e. by listing all its possible variations (given the range of representations employed in the protocols) would be unnecessarily unwieldy. A simpler and more fruitful approach, which has been adopted here, is to define it by *intension*, i.e. by an overarching principle instead. How this principle was arrived at is described below.

Some design actions are easily differentiable. In the earlier phases of the protocols they often consist of an individual physical act, such as making a mark on paper with a single stroke:

Begin End

Figure 3-5 an individual design action

It can be seen in Figure 3-5 that, as well as consisting of easily identifiable and discrete physical action, each of these earlier design actions has a definite endpoint, namely the point at which a designer pauses to assess the effects of their preceding design action on the representation. Each design action, therefore, is clearly differentiable from its neighbours by a cessation in the designer’s *modification* of the design representation.

This cycle of action and assessment, which was observed throughout all the protocols, operates on a finer grain scale than perhaps Simon envisaged in his generate-test model of the design process but would still appear to be commensurate with it. The relationship between this model and Schön and Wiggin’s see-move-see model, previously highlighted in the literature review, now serves to define those design actions which are not composed of single, easily differentiable physical actions. The observation made then, that Schön and Wiggin’s model could be viewed in the light of Simon’s model as *test*-generate-test, can now be used to mark a design action’s beginning as well as its end. If the end of a design action (generate) is marked by a pause for assessment (test), then its commencement is marked by the pause for assessment (test) of the preceding design action. A single design action is defined therefore as any form of modification of a design representation that is bounded by these two pauses.

Where compound actions occur, such as rehearsing the stroke of a pencil before committing a mark to paper, repeatedly tracing over the same line to reinforce it, or applying a number of sanding strokes to a block of foam, these have been concatenated into a single design action. In each of these cases to regard an individual stroke of the pen, or stroke of a sheet of glass-paper, as a discrete design action would be to invite even more unwieldy descriptions than those presented in the following chapters. The result would be to fragment the descriptions of design activity to the point of unintelligibility, while not necessarily being any more representative of a designer's intentions.

Individual design actions are then collected into episodes. When design actions are carried out on a two-dimensional representation an episode is defined as commencing with the first design action in that representation and terminating with the first design action carried out in a different type, or different instance, of design representation. In the following chapters timelines showing the general extent of these episodes appear at the head of the relevant section. The section heading itself states the kind of design representation employed, as well as giving specific timings for the beginning and end of each episode in minutes and seconds elapsed since the beginning of the protocol.

While the initial definition of an episode sufficed when participants were working in two-dimensional representations, during the latter stages of those protocols where they went on to make a three-dimensional representation, it was less obviously applicable. In these later episodes participants would often work on a single instance of a representation for considerable periods of time. It is possible to infer however that, although they were working on the same instance of a design representation throughout a particular period, its function as a design representation during this period could be divided into more or less distinct episodes. For example, participants' would often begin working on a three-dimensional representation as if it were an extrusion into three-dimensional space of the previously drawn two-dimensional shapes.

The key to dividing participants' actions on three-dimensional representations into episodes therefore lies in disentangling the spatial degree of the representation itself from the spatial degree of the design actions carried out on it. This is done by mapping the spatial degree of a design element acted on in a design action, and the degree of the space that element is transformed in, against Stiny's (2001) algebras of shapes:

$$\begin{array}{cccc} U_{00} & U_{01} & U_{02} & U_{03} \\ & U_{11} & U_{12} & U_{13} \\ & & U_{22} & U_{23} \\ & & & U_{33} \end{array}$$

Figure 3-6 the array of algebras of shapes

Algebras of shapes, $\boldsymbol{U_{ij}}$, are particularly appropriate for this task as they allow a clear distinction to be drawn between the dimension of the design element itself, ***i***, and the dimension of the space, ***j***, it is transformed in. Drawing for example, as it employs elements such as points, lines and planes which are transformed in the two-dimensional space of a sheet of paper, is contained in algebras $\boldsymbol{U_{02}}$, $\boldsymbol{U_{12}}$ and $\boldsymbol{U_{22}}$.

3.4.2 Describing design activity

The shape of these elements has been recreated in CAD on an action by action basis. Their fidelity has been guaranteed as much as possible by working directly over scans of the participants' drawings. These transitory shapes are extracted from particular actions either by observing them directly in the video record as they are created, or by examining the video record for individual drawing strokes (or other making actions) which match the shapes still present in the final representation:

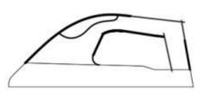

Figure 3-7 a CAD recreation of a transitory state

The sequential modifications of the designer's representation, observed in individual design actions extracted from the video record, are then applied in the same sequence to the CAD representation. The result is presented as a series of productions:

Figure 3-8 modifications are presented as production systems

While the depictions of the design representation shown in Figure 3-8 accurately model the initial and subsequent states, they suggest that all of the representation is required for the production to act on it. An alternative approach would be to abstract the modification required to transform the representation from its initial state to its subsequent one in the form of a shape rule:

Figure 3-9 the abstracted modification expressed as a shape rule

Shape rules are, however, perhaps too abstract when employed on their own; spatial relations are not easily read off from them, even in a strictly two-dimensional, linear representation. When the enquiry moves on to describing participant's design actions on three-dimensional representations, with their associated rotations and potentially infinite changes of view, depicting design actions in this fashion would be problematic.

While the depictions in Figure 3-8 explicitly capture the changing state of the representation, and the shape rule in Figure 3-9 captures the processes required to effect those changes, each was in its own way inadequate when it came to presenting the results of the protocols. To overcome these shortcomings a notation based on the physical analogue, or predecessor, of shape grammars, based on the use of yellow trace noted in the literature review, has been adopted here.

Shape rules themselves are of course at the heart of shape grammars. The way in which shape grammars employ Euclidean transformations, and sub-shape detection, to match the shape on the right hand side of a Post production system with an element of a target shape is analogous to a practice employed by architects in the pre-digital era. As was noted in the literature review, this entails physically moving shapes drawn on a sheet of yellow trace over the top of a design sketch until a satisfactory correlation is noted. When a shape grammar employs the Boolean operators of sum, union and difference to generate a new, third shape from the existing two, it emulates the action of an architect laying a fresh sheet of yellow trace on top of the previous one and drawing the derived shape on it. This action has been simulated here to establish the spatial relation of the shape rule to the representation as a whole, whilst emphasizing that its application is not dependent on the whole of the representation:

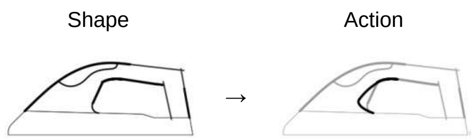

Figure 3-10 the 'yellow trace' notation finally adopted

A shape grammar, as an algorithm for calculating with shape (Stiny 2006) also needs unambiguously defined start and end points. In shape grammars these are supplied by initial and terminal shapes. The sequences of actions contained in the individual episodes, while employing shape rules in a similar fashion to a shape grammar, instead remain open-ended.

Although each episode has an initial shape, switching to another representation requires either a new initial shape or, when a previously employed representation is returned to, an existing one in the form of the product of the last shape rule application in that representation. In both cases this shape is treated (for the purposes of laying out the data in the following chapters) as if it were an initial shape in a shape grammar, and therefore as the start of a new episode.

Again, to reflect the open-ended nature of the sequences of actions contained in episodes, what would normally require a terminal shape in a shape grammar is instead replaced by a 'resultant' shape here. This shape is the final one of the grammar but, rather than terminating the grammar definitively it merely suspends it indefinitely.

So, while it may be clear that a new initial shape entails a new grammar, it could be argued that returning to the product of the most recent shape rule in a previously employed representation is simply a continuation of that earlier grammar. Although the resultant shape of the previous grammar may be geometrically and topologically identical to that employed as the initial shape of a purportedly new grammar, the designer's perception of it will have altered after an intervening period working on a different set of shapes in the preceding representation. It can be argued that, as such, it represents a different shape to the designer than that generated by the earlier grammar. Therefore, when laying out the results of the protocols, initial and resultant shapes are illustrated where a change of representation has been made:

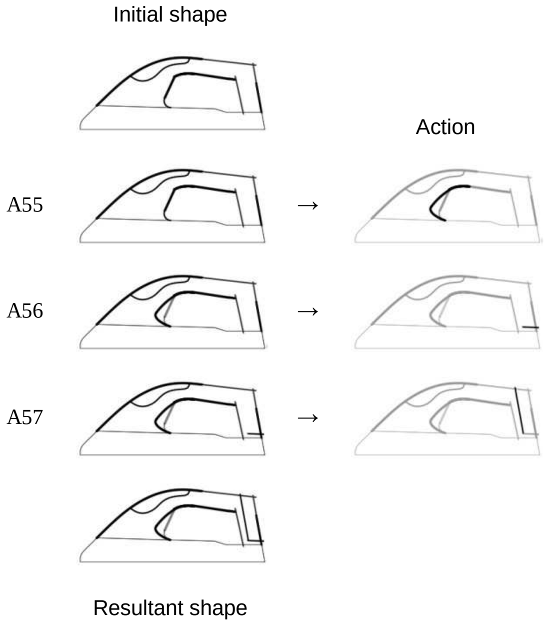

Figure 3-11 tables of design actions

Finally, these abstracted shape rules are matched up with the pair of images that define the beginning and end of the *physical* events they were derived from:

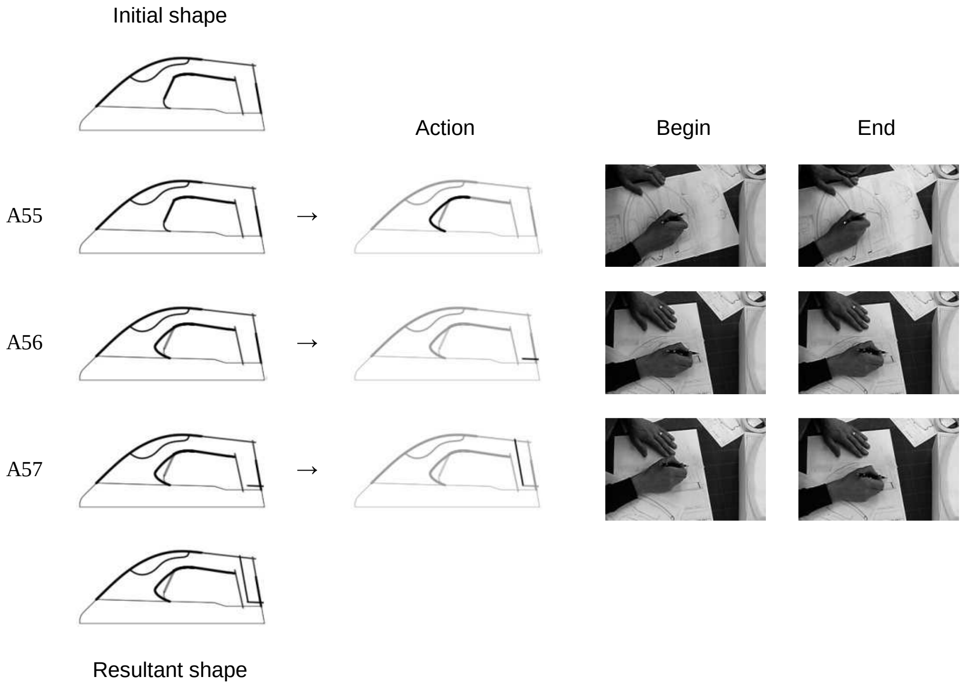

Figure 3-12 linked with corresponding physical activity

Linking the video stills showing the beginning and end of each design action to each shape grammar rule, as well as helping to make the varying definition of what constitutes the design action underlying the rule clearer in each case, also acts on a more superficial level as a reminder of the very human, embodied and situated nature of design activity.

More importantly perhaps, and an unforeseen consequence of this insistence on providing visual evidence of the beginning and ending of each design action, by recreating them in CAD, has been to provide a test of the validity of the rules identified by the author and the validity of the sequence of their application. Although a great deal of care was taken to identify and sequence rules thoroughly and correctly, and the initial laying out in tabular form looked convincing as a true representation of what had been observed taking place in the video record of the protocols, it became clear when first attempting to link these rules to stills taken from the video record of the beginning and end of the design action that they described, that steps had been missed or placed out of sequence.

The tendency to impose a logical flow to events retrospectively is powerful and insidious, a natural consequence of our human need to make an easily assimilable narrative out of events that may happen in an apparently illogical sequence, or that simply happen too quickly to be followed by ordinary human senses. Matching the CAD recreations of design actions against images of their beginning and endpoints taken from the video record on a one-by-one basis, although laborious, identified areas where false assumptions had been made in identifying rules, and in sequencing them, which might have escaped attention if they had been based on viewing the video alone.

3.5 Endnotes

Although verbal descriptions are widespread in protocol analysis, and have been applied successfully in other areas of research which deal with declarative knowledge, this enquiry is concerned instead with the richer world of shape. As it seems unlikely that designers' verbalisations can add to enquiries such as this, which necessarily deal with non-verbal knowledge and reasoning, a new approach to accessing this knowledge is required. The guiding principle underlying all the data analysis methods described in this chapter, as it concentrates on the interactions between designers and their representations, is that actions speak louder than words. The particular actions focused on are those which designers undertake in the process of creating their geometrical design representations. It is the sequence in which these actions, and interactions, take place that was ultimately analysed as a key to understanding how that process generates three-dimensional form.

By modelling architects' use of yellow trace as Post production systems, shape grammars have successfully been used to capture the formal essence of architectural plan design. The shape grammar formalism has been employed here in the opposite sense however. Rather than looking at a finished object and extracting a possible set of rules to generate that object from, this enquiry instead looks at how representations change and represents that change as sets of shape rules. These shape rules allow the enquiry to get inside the *process* of representation making, rather than merely viewing the end products of that representation making process and trying to infer rules from it.

The same sets of rules that have been derived from this examination of the representation making process can, of course, also be used to generate wholly original designs as well as to recreate existing ones. The shape rules, as presented in the following pages, are however decomposed more than is strictly necessary for a shape grammar. If the author's intention had solely been to produce a grammar that would recreate the objects produced by the participants in these protocols then many rules could be usefully combined. However, as the purpose of the protocols is to capture the complexity of the interaction between designers and their design representations, and to track that interaction across an indefinite number and variety of representations, it has been necessary to enumerate shape rules on an action-by-action basis. The more familiar linear treatment of shape grammars has also had to be extended in the representations of shape used in many of the design rules. Firstly to suggest participant's use of different line-weights, which often increased throughout the protocol to suggest a growing commitment to one particular line out of a number of options, as well as encompassing other non-linear marks such as those employed in the shaded areas.

In total, seven design protocols were recorded for this phase of the enquiry. The nature of the representations employed by participants in their protocols ranged from purely two-dimensional physical drawings to fully three-dimensional virtual geometry. Four produced blue foam physical representations, as well as side elevation and plan view drawings, templates and various supplementary drawings. One sculpted a half-scale clay representation, in addition to side elevation and plan view drawings and supplementary drawings. One produced virtual three-dimensional geometry in addition to side elevation drawings and a printed side elevation and one produced side elevation and plan view drawings alone without any form of three-dimensional representation.

Four of these protocols are presented in the following chapters, all of which were undertaken by experienced designers. The reason for focusing exclusively on the experienced designers is that their protocols are more likely to be representative of professional design practice and, furthermore, to avoid any unnecessary redundancy in what will necessarily be a finely detailed examination of the remaining participants' activities.

The protocols themselves are presented here in order of increasingly higher spatial degree of the final representation employed. They progress from strictly two-dimensional, to two-and-a-half-dimensional, to fully three-dimensional. The final protocol also results in a three-dimensional representation but approaches making that three-dimensional representation through a virtual instead of a physical medium.