KUNGL TEKNISKA HÖGSKOLAN
Royal Institute of Technology
Numerical Analysis and Computing Science

TRITA-NA-D9702 • CID-9, KTH, Stockholm, Sweden 1997

Collaborative Virtual Environments:

Grounding Development and Evaluation in Social Scientific Analysis

John Bowers

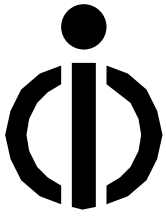

CID
Centre for
User Oriented IT Design

John Bowers

Collaborative Virtual Environments:
Grounding Development and Evaluation in Social Scientific Analysis

Report number: TRITA-NA-D9702, CID-9

Publication date: August 1997

E-mail of author: bowers@hera.psy.man.ac.uk

Reports can be ordered from:

CID, Centre for User Oriented IT Design
Nada, Dept. Computing Science
KTH, Royal Institute of Technology
S-100 44 Stockhom, Sweden
telephone: + 46 8 790 91 00
fax: + 46 8 790 90 99
e-mail: cid@nada.kth.se
URL: http://www.nada.kth.se/cid/

Collaborative Virtual Environments:

Grounding Design and Evaluation in Social Scientific Analysis

A compilation of three papers edited and introduced by John Bowers

General Introduction

This report comprises three papers documenting recent work of relevance to collaborative virtual environments (CVEs) by John Bowers (of CID, and the Department of Psychology at the University of Manchester, UK), James Pycock (formerly of the Department of Psychology, University of Manchester, UK, now at Rank Xerox Research Laboratory, Cambridge, UK), and in collaboration with Jon O'Brien of the Sociology Department, Lancaster University (UK). While these papers have been published before separately, this report gathers them together to make them more accessible as a collection of related studies and to enable comparisons to be made of their research methods and the different ways in which they attempt to ground system development and evaluation in social scientific study.

The first paper presents some analyses of interaction in a virtual meeting held within a CVE using social scientific methods employing talk and (virtual) body movement transcription to investigate how users deploy their communicative competencies in such new environments.

The second paper complements these methods with ethnographic observation of a series of meetings in CVEs to develop the argument that CVEs can only be fully understood and evaluated if one examines social interaction within both the virtual and real worlds (both

the common electronic environment that participants share and the physical environments at their respective sites) and the relation between these two.

The final paper—rather than looking at social interaction in already existing CVEs—looks at the question of how requirements for future CVEs can be developed using ethnographic research techniques. Our research has focused on the potential role of CVEs (and generally of VR and 3D visualisation) to support cooperative work in the fashion industry, especially the work of fashion designers.

The three papers—taken together—display the variety of research methods we have employed (interaction analysis and ethnography) and our interest in exploring the use of social scientific research at several moments in the 'life cycle' of development of CVEs (both at 'requirements' and in 'evaluation'). Along the way, we believe we have developed a novel social scientifically informed approach both to the analysis of activity in CVEs and to certain foundational concepts in VR research (e.g. 'immersion').

John Bowers, February 1997

Talk and Embodiment in Collaborative Virtual Environments

John Bowers, James Pycock and Jon O'Brien

Abstract

This paper presents some qualitative, interpretative analyses of social interaction in an internationally distributed, real-time, multi-party meeting held within a collaborative virtual environment (CVE). The analyses reveal some systematic problems with turn taking and participation in such environments. We also examine how the simple polygonal shapes by means of which users were represented and embodied in the environment are deployed in social interaction. Strikingly, some familiar coordinations of body movement are observed even though such embodiments are very minimal shapes. The paper concludes with some suggestions for technical development, derived from the empirical analyses, which might enhance interactivity in virtual worlds for collaboration and cooperative work.

Introduction

Collaborative Virtual Environments (CVEs), where multiple individuals interact with each other in a computational environment rendered by Virtual Reality (VR) technology, are of emerging interest from a number of perspectives. Several writers have claimed [e.g. 2] that such environments may support collaboration and interactivity (especially geographically distributed collaboration) in ways which go beyond what is possible using more familiar meeting room or teleconferencing technologies. The development of CVEs has also drawn interest as a complementary alternative to video and 'mediaspace' [5, 12] research. CVEs may provide a shared spatial environment where (in principle) people can employ communicative resources which are unavailable to them in other technical systems. For example, participants can have a degree of control over what they view in a CVE which is not generally possible with mediaspaces supported by a fixed camera and monitor system, and our ordinary means for coordinating turn-taking in social interaction can be deployed rather than some technical means such as 'floor control' policies as is commonly the case in traditional conferencing systems. Furthermore, if users in a CVE are all embodied in it so that their location and orientation can be represented, then a degree of mutual awareness of each other's activity may arise or be readily supported [1, 2].

This paper attempts to subject such claims to preliminary examination by giving a characterisation of what CVEs are like as environments for cooperative work and social interaction and seeing how ordinary conversational mechanisms are exploited or transformed in such environments. To these ends we employ empirical techniques derived from Conversation Analysis or CA [e.g. 14], which to our knowledge has not been attempted before in studying CVEs or other VR technologies, though CA has had some influence in HCI and CSCW in enabling detailed studies of interaction in the workplace [10], the impacts of new computer based technologies on 'talk at work' [8] as well as in motivating technical design choices [3] and assisting in the analysis of the design process itself [4]. Accordingly, we seek to add to this literature while extending it to the study of a novel setting (a work-related meeting being conducted in VR).

Finally, we are concerned to show how methods of interaction analysis might contribute to the evaluation and hence future requirements of CVEs. User-oriented evaluative studies of VR systems are still overwhelmingly dominated by investigations of such matters as motion sickness [13] and the characterisation of phenomena in terms of individual perceptual psychology [15]. We wish to study CVEs so as to extend the base in reference to

which VR systems should be evaluated and developed. Through characterising the nature of social interaction in a CVE at least in a preliminary fashion, we hope to make a start on the task of exploring the worth of distributed VR technologies for the support of cooperative work and social interaction.

The Virtual Meeting

We have been working closely with CVE developers for some time in a number of research projects. As their CVE technologies have become more robust, our collaborators have begun to use their own systems as environments for the support of distributed research meetings. We have been present as participant-observers at a number of these meetings and have begun to accumulate a corpus of videos and other materials for analysis. At the time of its occurrence, probably the most significant of these meetings took place on 28th March 1995 as a working meeting of COMIC, a European ESPRIT project devoted to basic research on CSCW. The significance of this meeting was that it involved real-time interaction in a number of virtual worlds between nine participants who were distributed across five sites and three countries (Germany, Sweden and the UK). According to many of the participants, this degree of international participation was a 'world first' and has been reported as such in a number of technical press articles.

The MASSIVE system [9] that was used for the virtual meeting supports multi-user interaction between distributed sites allowing participants to communicate over graphical, textual and audio media. The graphical interface provides a navigable 3D view of the shared virtual world and of other participants represented as simple graphical embodiments. For the current meeting the 3D view was presented on screen rather than immersively. The audio interface allows real-time conversation. The text interface provides a 2D plan view of a world and allows the exchange of text messages. MASSIVE employs client-server information distribution. The hardware used was exclusively Silicon Graphics for the interface clients with a Sun Sparc 10/51 at Nottingham University in the UK running server software.

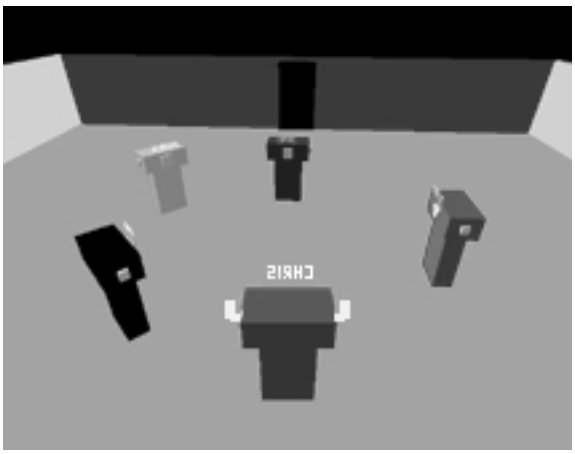

Figure 1. Blockies in a virtual meeting.

The user-embodiments or 'blockies' are made of simple 3D box-polygons with one square 'eye' on one vertical surface and the user's name 'suspended' above the top surface. As well as identity, this design affords a rudimentary sense of 'face', 'front' and 'back' which according to Goffman [6] are features of the human body of basic significance to social interaction enabling us to distinguish between, for example, talking to someone's face from talking behind their backs. Furthermore, the 3D view that a user has can relate to the blockie's face. Although different views are possible, the default is that one 'sees out of' the blockie's eye. This view, or one where one looks over the blockie's shoul-

ders, are the ones typically employed by users. In this way, what other participants can see and where they are looking is often available from an inspection of their bodily orientation and, accordingly, a sense of mutual awareness can be sustained and transformed by aligning the blockies or moving them around. For example, 'full face encounters' [6] can be brought about by two participants aligning their blockies to face each other. Indeed, it has been precisely a consideration of what was the minimal geometrical object necessary to sustain basic interactional relations between participants which has informed the design of the blockies [1]. Insisting that the embodiments should be geometrically simple (yet still have interactional potential) is necessary because of the extreme computational complexity of distributed VR systems based on current technology.

The blockies also support minimal gesturing. They have 'ears' which can be 'flapped' in different ways (left one raised, right raised etc.). They can recline (or 'sleep'). This can be used, for example, to denote that the participant the blockie corresponds to has currently left their local machine and is not available for interaction. Gestures are controlled through simple key sequences and the blockie is moved by clicking the mouse on the 3D view or by using the arrow keys. Finally, the blockies have a 'mouth' which opens when a user's speech exceeds a certain amplitude threshold.

The overall business of the virtual meeting is well described by the following quote from the minutes which one of the participants distributed afterwards: "The time spent in the conferencing software, approximately an hour and a half, can be split up into three distinct periods. To begin with there was a fifteen minute mingling session where people arrived for the meeting and chatted socially ... This time was also spent making sure everybody could communicate with everybody else. At a quarter past two SB called the meeting to order and everybody trooped off to the designated meeting room where a pre-prepared agenda was awaiting on a noticeboard and ... introductions were carried out. CG then gave a tutorial on how to use MASSIVE ... [The items on the agenda were then discussed.] The formal meeting finished at about two fifty five, and after that much more informal communication took place." While some of the clients needed to be restarted from time to time, the server software and network connections were reliable throughout.

Interaction Analysis

In this section, we report our analyses of the talk exchanged between participants in the virtual meeting. Specifically, we focus on two issues: the nature of turn taking in virtual environments and how the embodiments are used. Our interest is in the *interactional qualities* of CVEs, what these environments are like as arenas for social interaction as revealed in the moment-by-moment texture of talk and activity. Accordingly, our analyses work from transcript data which we analyse qualitatively and interpretatively.

Before we turn to some examples from our data, it is necessary to clarify the transcript conventions we have adopted. We employ an adapted version of the conventions devised by Jefferson and presented in a number of sources [e.g. 14]. Pauses and silences are notated by their length in seconds shown within round brackets: `(1.2)`. Talk which receives more emphasis than the surrounding speech is underlined: `someone else's turn`. Parts of the transcript where we are unsure of what is said, but are able to guess, are notated with round brackets: `thanks lennart (how) very eloquent`. Where we are unable to guess the round brackets are empty. Prolonged sounds are indicated by inserting colons, and concatenated speech, where the words are quickly run together, is notated by hyphens placed between the words: `an::d this wonderful VR system-MASSIVE-that-we're-using is wot i writ`. An audible in-breath is notated `.hhh` and `hhh.` denotes an audible out-breath. Overlapping speech is notated by means of square brackets positioned to indicate where the overlaps occur. Speech which is distorted by, say, some malfunctioning of the audio link is placed within curly brackets: `i'm john {b}owers from manchester {uni}versity`.

Our comments within the transcript are enclosed within angled brackets: <laughter>. We shall discuss our conventions for transcribing the movements of the blockies later.

Taking Turns at Talk

Turn taking is a basic and obvious feature of the organization of talk. Some aspects of how turn taking is managed by participants in ordinary conversation are analysed by Sacks, Schegloff and Jefferson [14] who propose what they call a 'simplest systematics' by means of which participants manage the exchange of speakers from turn to turn. In the analysis of turn taking, a basic distinction can be made between (i) turns which select who is next to speak by, for example, addressing a question (ahh lennart can you hear us?) or request (go on dave) to a specific, perhaps explicitly named, participant and (ii) turns which do not contain next-selecting components:

```
AB:   <mouth click> .hhh (2.2) <mouth click> .hhh-woahruh i'll go next
      (.) then if no one else is speaking. (0.6) uh i'm adrian bullock
      also from the university of nottingham.
```

In AB's turn here, it is AB who selects himself to speak, indeed he does so quite explicitly noting that 'self-selection' is exactly what he is doing (i'll go next). AB's turn itself contains no components which project who is next to speak after him. He introduces himself and then stops. At such moments, it is open for the other participants to either select themselves as next-to-speak or, for that matter, for AB to continue speaking prolonging his current turn.

Next Selected by Self

Our preliminary observations suggested to us that turn taking was often problematic in the virtual meeting. The transcript was marked by many long silences as participants seemed to wait for each other to say things. However, a closer inspection reveals that these silences are most prevalent when speakers have to select themselves as next-to-speak in the absence of any prior turn with next-selecting components. This is not to say that all such turn transitions break down. Rather it is to claim that the disfluencies in speech exchange are notably concentrated at transitions where speakers have to self-select. The following examples (from the early stages of the meeting where SB has invited participants to introduce themselves) reveal problems in speaker switches of this sort.

Example 1

```
SB:   someone else's turn. (1.2) thanks lennart (how) very eloquent.
      (1.2) maybe (eloquent.)
(12.0)
CG:   i'll 'ave a go then.
(1.2)
SB:   yeah please do.
(1.4)
CG:   i'm chris greenhalgh (0.4) also at nottingham (0.8) an::d this
      wonderful VR system-MASSIVE-that-we're-using is wot i writ.
```

Example 2

```
SB:   <laughter> how much does it cost?
CG:   oo to you nothing. bargain at double the price.
(1.0)
SB:   excellent.
(3.0)
```

```
AB:   <mouth click> .hhh (2.2) <mouth click> .hhh-woahruh i'll go next
      (.) then if no one else is speaking. (0.6) uh i'm adrian bullock
      also from the university of nottingham.
```

Example 3

```
SB:   anyone else?
(1.2)
AC:   {uh uh hello s-s-} can you hear me?
      [ (.)  ]=anyone? [ (.) ]=err
SB:   [ yeah=]         [ yea=]
AC:   it's andy at lancaster.
```

Example 4

```
(1.2)
( ):  (.hhh)
(4.0)
SB:   anyone else?
(2.8)
JB:   huhlow::. i'm john {b}owers from manchester {uni}versity cur-
      rently v{i}sit{in}g sics. (0.8) i'm interes{ted in} lots-of
      stuff (0.4) {and i'm a c}apricorn.
SB:   <laughter>
```

In Examples 1, 3 and 4, SB does not explicitly name who is next to introduce themselves. Rather he invites contributions by means of `anyone else?` (Examples 3 and 4) or `someone else's turn` (Example 1), leaving it up to whoever is next to speak to select themselves. Such components are followed by quite lengthy silences. Indeed, in Example 1, after just over a second's silence, SB engages in some 'side talk' thanking LF in an ironically humorous fashion for his prior self-introduction. Another twelve seconds of silence follows before CG starts talking. In Example 2, SB and CG exchange some humorous remarks before SB utters `excellent` which can be heard as closing his exchange with CG. The three second pause that follows is then broken by AB selecting himself. In all of these examples, speakers only self-select after quite lengthy silences and with much preparatory activity (e.g. mouth clicks, in-breaths, protracted sounds, stammerings and so forth) or, as in Example 1, a minimal turn from CG (`i'll 'ave a go then`) which requires confirmation from SB (`yeah please do`) before CG continues.

This preparatory activity is often quite exaggerated as in Example 2 where a mouth click and an audible in-breath are heard from AB and then, after a pause of just over two seconds, there follows a second mouth click, another audible in-breath running into a vocalization transcribed as `woahruh` before AB explicitly self-selects. The interactional significance of such preparatory activity is worth noting. Audible in-breaths and the rest do not explicitly or fully claim a turn at talk in and of themselves. They could be interrupted by another participant immediately launching into a turn without such activity. Accordingly, such preparatory activity displays a participant's *readiness* to contribute as next-to-speak, without disqualifying others. Indeed, in Example 2, even *after* he self-selects, AB pauses very briefly (notated by (.)) after `i'll go next` and for about six tenths of a second after `then if no one else is speaking`. These are further junctures where another participant could have self-selected and claimed the floor to introduce themselves ahead of AB. These features of the examples suggest that self-selected turns are managed with considerable care by speakers - a matter borne out by the fact that AB in Example 2 is explicitly attentive to the possibility that others may speak ahead of him. The exaggeration of preparatory components in the virtual meeting is, we suggest, a means for managing turn taking at moments which can be problematic where, for example, in the absence of explicit

next-selection, a number of speakers could start to speak simultaneously. Indeed, Examples 5 and 6 suggest that simultaneous self-selected turns are problematic for the smooth conduct of the virtual meeting and that the presence of audio distortion makes them especially hard to manage.

Example 5

```
(3.0)
SB:   [ (  )
JB:   [ (      )(fahlén)
(3.7)
LF:   okay. errm. (2.5) we have err (.) umm (.) a rather a serious
      echo here.
```

Example 6

```
(1.8)
LF:   {i think we ermm} (0.5)
KJ:   [ {(let's put) that upon} the agenda ] (1.0)
LF:   [ {(                              )} ] {uh }
KJ:   [ (     ) ]
LF:   [ {(      ] ) some kind of uh distortion (.) uhhhm (0.5) origi-
      nating at K T H. you should probably lower yer input levels a
      bit.}
(1.0)
KJ:   okay i'll try.
```

In Example 5, after a three second silence, SB and JB start speaking together. JB's speech continues after SB's but neither of their turns are taken up by any other participant, rather the talk dies away into a silence of nearly four seconds before LF complains about the audio quality. In Example 6, KJ and LF talk together and then after LF utters `uh` while KJ falls silent, they both restart together again with LF's complaints about the audio again heard after KJ's talk finishes. In cases like these, the overlaps leave unfinished interactional business with one or both overlapping turns failing to be taken up by others in subsequent talk. The overlaps in Examples 5 and 6 are hard to transcribe for the precise reason that the voices tend to mask each other and this is exacerbated by the poor audio quality LF complains about. Not only do they mask each other so as to make transcription difficult, they interfere so as to make it hard for participants themselves to hear the talk and tease apart overlapping contributions. We suggest that it is because of the problems involved in the management of overlaps of turns that self-selected turns often manifest the preparatory activity we have noted.

Note again the presence of various artifacts (e.g. distortion) in JB's talk in Example 4. In the next example, this is particularly intense. SB does not hear DE's `eh:::m:. (0.5) yuh` as preparatory to a turn at all!

Example 7

```
(2.0)
SB:   okay-is there any shy: person who hasn't spoken yet?
(3.0)
DE:   {eh:::m:.} (1.5) {yuh}
SB:   i think that's everyone.=
DE:   ={ouh}=
LF:   =steve-d'yuh-think-i think you could probably
      [ errr (0.9) ] raise your levels a bit?
```

```
DE:   [ {hallo::?} ]
SB:   who me? maybe.=
DE:   ={huhlo?}
AB:   <mouth click> yeah go on dave.
DE:   (2.5) {right. um. [ (.)  ] (first)ly can you hear me:?}=
LF:                     [ (  ) ]
      =you're much quieter [ (  )
AB:                        [ just about.
(2.8)
SB:   hi the[re.
AB:         [{ } breaking up (.) but you can just about make it out.
(0.8)
( ):  um.
(2.0)
DE:   {{(okay)}} (1.6) {{{(hiya (.) steve)(   [
      [     [          )}}}
JB:                                          [<laughter>
SB:   [ ah ha. um.
AB:         [nah.
AB:   nah-that's unintelligible.
```

This example suggests the problems in coordinating speaker switches when at least one participant has a considerably degraded signal may be profound. It also further underlines the point we have been making about difficulties with self-selection and that these may be intensified when one's signal is poor. Indeed, very few of DE's turns at talk are brought off in Example 7 without some overlap or interruption from the other participants.

Interestingly, DE's first attempt to claim a self-selected turn at talk in this example manifests very similar preparatory features to those we have already seen in Examples 1 to 4. It is unsuccessful presumably because SB does not hear them as having this significance or as being any different from the artifacts, pops and crackles and other background noises that can be heard on the audio channel. DE having the poorest audio connection is doubly disadvantaged: first in that his speech is easily masked by others in overlap, secondly in that his routine attempts to anticipate the problems of overlap (e.g. protracting an ehm or uttering a preparatory yup) are not heard as such!

It is important to emphasise that the difficulties with turn taking we have noted cannot simply be reduced to problems with audio quality. We observe substantial silences in examples where there are no (Examples 1 and 2) or few (Examples 3 and 4) audio problems. Indeed the audio connections within the site accommodating SB, CG and AB were of good quality throughout the meeting, yet these provide some of the most notable silences in our transcripts. Hence we argue that the problems of self-selection are exacerbated by but not solely attributable to poor audio quality.

The lengthy silences before speaker-switches, we suggest, reflect problems due to managing self-selection with minimal embodiments which have restricted gestural abilities. In this regard, it is of interest that DE does not attempt to compensate for interactional difficulties by any form of virtual gesture or change of body orientation at any moment in Example 7. The embodiments are very rarely used concurrently to aid speakers in designing their own turns or in eliciting turns from others at such moments (in contrast to the use of gesture and body movement in ordinary co-present conversation, see [7, 11], or, for that matter, as reported in the videoconferencing literature, see [16]). This is a point we shall return to.

Next Selected by Prior Speaker

In contrast to the problems we have noted at moments where speakers have to self-select to claim a turn at talk, turn transitions are much smoother where the identity of the next person to speak is clear either because they have been explicitly named or due to some other contextual feature. In Example 8, AB gives JB instruction as to how to use the mouse button after he has been struggling to pass through a gateway to another virtual world. Similarly, in Example 9 KJ and SB discuss two different modes for 'focus' a technical feature of the MASSIVE system. Finally, in Example 10 SB notices the presence of LF and starts a discussion about the organization of a conference panel.

Example 8

```
AB:   (you can use your right mouse button for lateral) translations
      john.
JB:   i can do what?
AB:   use your right mouse button (.) in the middle of the screen just
      to move you side ways.
<JB moves sideways>
AB:   wrong way.
SB:   oh he's gone. <through the gateway>
```

Example 9

```
(KJ): (  )
SB:   (did you) change focus?
(0.5)
KJ:   so what is the difference between narrow and directed
SB:   directed is a sort of cone (.) in front of you erm i can't re-
      member the angle it goes out to. narrow is like a really thin
      tube
(0.5)
KJ:   ok
(0.5)
SB:   you can try this thing out in the audio gallery and (you'll)
      probably notice the effect fairly well
KJ:   alright
```

Example 10

```
SB:   oh there you are. ok so w-what happened about this panel thing
      what's the (scam)?
LF:   umm nothing much but erm we're still missing erm erm: the erm er
      spatial awareness versus {(   )} versus media spaces kind of pa-
      nel ya?
SB:   arham are we liable t-to do one of these?
LF:   ya <conversation between SB and LF continues for nearly a minute
      with many turns readily exchanged>
```

In all these examples, the speech exchange between participants is fluent and not marked by the hesitations, overlaps and other perturbations noted before. In each of them, though, who is next to speak is clear either as a result of explicit naming or the identification of 'you' even if several other parties are present. When speaker switches can be governed in this way, fluent conversations can occur between two parties in the virtual environment. (It should be noted that these examples involve fluent exchanges between people at distant sites: SB and AB are in Nottingham, the others in Stockholm. This suggest

that the problems we have documented are interactional and not just due solely to technical reasons like network delays which were rarely notable in the meeting.)

Interpreting Silences and Absences

We have suggested that there are differences between turn transitions where the prior speaker has selected who is next to speak (these are carried off relatively fluently) and those where speakers have to self-select (these are the sites of the main breakdowns of fluency in the transcript). As noted by Sacks et al. [14], these different transition types give rise to differences in the significance of silences in talk. When the prior speaker has selected who is next-to-speak, any silence after next-selection is *attributable.*That is, it is 'owned by' the person who has been designated as next to speak. It is *their* silence and their next speech action is *noticeably absent.* Depending on the circumstances, a continued silence after next has been explicitly selected may licence inferences about the person who is remaining silent. For example, in a law court, persistent silence under cross examination may lead to suspicions of guilt or other culpability. On the other hand, when the prior speaker has not selected who is next-to-speak, such attributions are not standardly available and the silence is not uniquely 'owned'.

When conversations are technically mediated, technical failure can be inferred as a source of attributable silence. Indeed, in the data we have, technical failure (rather than some socially significant attribution like evasiveness or rudeness) is *invariably* first considered as accounting for an attributable silence. In Example 11, SB first asks a question about whether another meeting is intended. AB replies to this saying that a future meeting should use the DIVE system. SB continues by asking a question which is hearably directed at two of the current meeting's members who are also developers of DIVE. However, this receives no reply and after a one second silence, SB explicitly names (and aligns his embodiment in the virtual world so as to face) the two the question is addressed to. Again, a long silence follows and AB checks on LF's ability to hear. Another long silence: whereupon SB notices and brings to the attention of others that KJ has typed a message in the text window saying he cannot hear anything. The point of this example is that when next-to-speak has been explicitly selected and no reply is heard, this is interpreted here as arising from an inability to hear due to an audio failure.

Example 11

```
SB:   are we going to have another one of these?
AB:   we ought to have one with dive.
SB:   yer (0.5) what's-what's-what's the status of that stuff? (1.0)
      he says looking at lennart and kai.
(5.0)
AB:   ah:: lennart can you hear us?
(5.0)
SB:   (hi) if you look at the text window (.) kai says he's not hea-
      ring anything
```

Example 12

```
AB:   can you hear me kai? (2.5) hello?
CG:   he may be a corpse now. he's disappeared off my text (  as far
      as  )
```

Similarly, in Example 12 AB's repeated failure to elicit a reply from KJ is accounted for by CG. Amongst the VR research groups involved in this meeting 'corpse' is used to refer to an embodiment which lingers in the virtual world even though the local system which manages interaction with it (in this case the interface client running on KJ's ma-

chine) may have crashed or become disconnected. In both examples, noticeably absent 'seconds' in an 'adjacency pair' [14] like question-answer are first interpreted as indexing *technical problems and not accountable social behaviour* (e.g. rudeness etc) on the part of whoever has been selected to speak. In contrast, the silences in Examples 1 to 4 where SB did not select who was next to speak, are not uniquely attributable and hence, if anyone does happen to have audio connection or other technical problems, they will pass unnoticed and possibly for some time to follow. To put this point another way, silence can be hard to interpret and technical problems can be difficult to identify, even those originating from one's own local environment. A local failure can pass unnoticed for some time simply because an appropriate moment for the problem to be identified does not occur in the interaction. We shall return to the design implications of this in discussion at the end of this paper.

Embodiment, Talk and Movement

How are the embodiments used by participants in the virtual meeting and how are their movements coordinated with concurrent talk? A first observation is that participants do move the embodiments around to get a better view of those they wish to interact with or attend to. The provision of an audio channel does not *require* them to do so in order to speak to each other, but that we have examples of people navigating the embodiments in this way suggests that, in a simple sense, interactants are seeking to become 'face engaged' [6] when exchanging talk. Accordingly, the embodiments do seem to have a social *interactional* role and not merely a role in determining the view an individual has of the virtual world. Note how SB and LF establish contact in Example 13 before starting a lengthy exchange.

Example 13

```
SB:   (   ) lennart
LF:   yar (steve)
SB:   oh what happened about um the the panel for ecscw?
(LF): (   )
SB:   c-could you turn (the) mic up a bit?
LF:   (   )
SB:   (   ) where has he gone lennart?
(2.0)
LF:   ar where's steve?
SB:   hi i'm i'm here behind kai i think
LF:   ye
(4.0)
<LF and SB move around the environment and finally turn to face each
      other>
SB:   oh there you are. <continues as Example 10>
```

Do we have evidence from the recordings of the virtual meeting of any more subtle uses of virtual body movements to accompany ongoing interaction? To examine this, we first consider whether body movements exist in the virtual world as an accompaniment to one's own talk and then look at whether participants are able to coordinate their movements with those of others and their talk.

We remarked above that participants seem to rarely use virtual body movements to aid the design of their own turns, even when constructing turns by means of talk alone is problematic (as in Example 7). Indeed, it is rather rare in the virtual meetings we have studied for people to complement their own turns at talk with any concurrent movement of their embodiment. This may not strike one as surprising as talking down a potentially

troublesome audio channel may be difficult enough without having to engage in simultaneous mouse movements to get one's embodiment to move! However, it presents a stark contrast with ordinary talk where a whole array of body and facial movements, gazings and changes in overall deportment can accompany and aid the design of turns at talk [7, 11]. Example 14 transcribes the body movement from Example 1. Here, CG raises and lowers the ears on his embodiment. These gestures span his breaking of the long silence we have already noted and aid the construction of his self-selected turn. This is however the only example we have yet found of gesture being used to aid the design of a speaker's concurrent turn. (It also introduces our transcript conventions for showing movement. The beginning and end of the movement are shown underneath the concurrent talk and described in italics on a line after that.)

Example 14

```
<The virtual body movements in Example 1>
      ( 12.0 )
CG:         ¶-
CG:   raises both ears

CG:   i'll 'ave a go then.
CG:   ------¶
CG:   and lowers them
```

While we have few examples of participants designing their own turns with the aid of some movement from the virtual embodiment, mutually coordinated movements between participants and movements coordinated with the speech of others are more commonly found in our data. Consider Example 15 which transcribes the body movements of a number of participants while LF is introducing himself to the meeting. We transcribe on the spot turnings about with ∞----∞ the length and the position of the symbol corresponding to the analogous position in the talk above it. We transcribe translation movements (xyz-displacements) with ^-----^. A verbal description of each movement is given just below each movement-transcription. A period (.) at the beginning of a line is used to match up lines of transcribed movement where no body movement occurred with the corresponding line of talk.

It is notoriously difficult to adequately and clearly transcribe body movement [8, 11] but we hope our conventions will become clear as we now explicate this example. Early in the example, AC turns towards AB just while LF utters `fahlén`. He then stays facing AB for the rest of LF's turn. A 0.6 second pause follows, at the very start of which AC begins to turn back towards SB, continuing this movement over a brief `uh hum` from SB and stopping the movement when LF begins to hesitate (`er:`) in his next turn. For his part, reciprocally, AB turns towards AC, again beginning the movement at a hesitancy in LF's turn (the initial `uhm:`). Immediately once AC has finished turning back towards SB and started to move away from the group, AB also turns back towards SB, starting this movement during LF's `er:`. AB makes a further movement towards SB while (again) LF is hesitating and pausing with `er:m (1.2)`. SB then makes a movement towards LF which starts during a 0.6 second pause in LF's talk. Following the start of SB's movement, LF continues talking with `and um:: i've been involved in these things for a long time`. When LF begins to utter `for a long time`, SB reciprocates AB's slightly earlier turn towards SB, before finally returning towards LF, initiating this movement again during a one second pause in LF's talk.

Example 15

```
<SB is facing the others who are looking towards him. After some side-
      talk with LF, SB invites LF to introduce himself>
LF:   (1.6) um: (2.0) i'm um (.) lennart fahlén from the swedish in-
      stitute of computer science
AC:                                      ∞----∞
      .
AC:   turns to face AB

      ( 0.6 )
AC:   ∞------
AC:   turns to face SB
SB:   uh hum=
AC:   -------
AC:   continues turning
LF:          =uhm: (1.0) and er: in my spare time i try: to manage
      this (    ) virtual reality group in here-at er:m (1.2) sics as
      it's called (0.6) and um:: i've been involved in these things
      for a long time (1.0) so i'm starting to (        ) (1.0)
      that's very apparent i think
AC:   ----------------------∞ ^-----------------------------------
      -------------------------------------------------^ ∞------∞
      .
      .
      .
AC:   finishes turning to face SB then backs away from the rest of the
      group then turns towards AB
AB:             ∞--------∞       ∞--------------------------------∞
                                                          ∞--∞
      .
      .
      .
AB:   turns towards AC then returns to face SB then turns further to-
      wards the rest of the group
SB:   .
      .
                  ^------^
      ∞---------------------∞                              ∞---
      -------------------∞
SB:   moves towards LF then turns towards AB and then turns back to-
      wards LF
```

We suggest that this example on close examination compellingly shows *exchanges* of movement between first AC and AB and then AB and SB during LF's introduction of himself. To the extent that the blockies can exchange glances at each other, this is what they are doing here (even though for a blockie a 'glance' involves a whole body movement!). Thus, participants do seem to be able to use the blockies to enter into coordinated movements with one another. The blockies have an *interactional* significance and are not merely navigation devices. What is more these virtual body movements seem to be very precisely coordinated with LF's ongoing talk. It is in the pauses and hesitations in LF's talk that we see the majority of the movements initiated. Thus, strikingly, participants seem to be able to coordinate movements with each other while interleaving this with the talk of a 'third party'. It is also noticeable that SB's approach to LF initiated during a 0.6

second pause provokes LF to further add to his turn of self-introduction. In all these respects, we are witnessing residual though noticeable phenomena reminiscent of the coordination of action in ordinarily embodied conversation. In short, even if only to a limited extent, participants do seem to be able use the blockies in ways derivative from their ordinary interactional competencies.

Discussion

We owe it to Heath et al. [12] to have drawn our attention to the following quote from a 1972 lecture of Harvey Sacks: "The technical apparatus [Sacks had in mind the telephone] is, then, being made at home with the rest of our world. And that's a thing that's routinely being done, and it's the source for the failures of technocratic dreams that if only we introduced some fantastic new communication machine the world will be transformed. What happens is that the object is made at home in the world that has whatever organisation it already has". We feel that our analyses of talk in virtual environments clearly exemplify Sacks' intuition that the new and technologically unfamiliar is often made to be 'at home' with our familiar world. If ever there were technocratic dreams, they have been thoroughly invested in VR! But what we see is the ordinary apparatus of conversation and the social interactional coordination of body movements being moulded and adapted to what the virtual environment affords so that participants can carry on as best they can with their business at hand.

Thus, we see systematic ways in which participants try to resolve or anticipate turn-taking problems, the elementary coordination of body movements between participants, the coordination of movements with ongoing speech, the utilisation of the bodies to engage others and initiate talk, amongst other phenomena. While, of course, we have concentrated on data from just one virtual meeting and it must be acknowledged that more meetings and more examples are required to develop yet more convincing generalisations, it is fascinating that what we have noted so far in the virtual world is in some way familiar.

Though familiar, this is not to say that improvements should not be made to systems such as MASSIVE to enhance their abilities to support multi-party collaborative activity in a virtual environment. Our interaction analytic techniques have been able to highlight problems, some of which may be amenable to technical solution or assistance. Let us discuss three classes of possibilities.

(1) We noted that it is possible for many technical failures to pass unnoticed for some time, simply because those moments which make them clear (a sufferer of technical failures being selected as next-to-speak, yet not responding) may not have arisen in the conduct of the meeting. This suggests to us that CVEs should support *local troubleshooting* because it may be very hard to bring to the attention of *others* that one is experiencing a local failure. This has implications for the overall distributed architecture that a system might exploit. An architecture must be not only robust in the face of local failure but it must also support graceful distributed recovery from local failures. A local failure must be remediable *at that site* and not require initiation from a remote site where participants may be unaware of the failure. Additionally, it should be possible to bring about such recovery by means within the expertise of any participant. These are actually quite demanding requirements and we believe they follow from our observations of how the structure of technically-mediated social interaction may lead to problems identifying ongoing failures in distributed systems such as MASSIVE.

(2) The overall design of virtual worlds should be considered in terms of how they afford social interaction and not just in terms, say, of their navigability, capability for presenting masses of information, or their thrilling aesthetics. The kinds of objects that we insert into a virtual world should be selected and designed with social interaction in mind. For example, a meeting table may be a simple device for people to gather around while affording them means for coordinating their talk, views of each other and mutual bodily orien-

tation. Indeed, such a device may have aided our participants in solving some (not all) of their turn-taking problems by suggesting a 'round the table' sequence for talk. Quite simple devices (e.g. a table as polygon on the base plane) may often be the most important from a social interactional standpoint yet their inclusion in the virtual world is easy to forget. Although the MASSIVE system can support a variety of 'meeting furniture', only a noticeboard was included on this occasion.

(3) While the blockies do afford various *social* interactional phenomena (and are not *merely* navigational aids or simple interface devices), it is worth reminding ourselves of the subtlety of interaction and participation which is much more readily possible with a real human body. The hands, the arms, the head, the neck, the torso permit a number of different orientations with respect to each other as well as with respect to co-interactants' bodies. In this way, we can glance without moving our heads, or turn our heads without moving the rest of our bodies.

Importantly, the coordinated flexibility of our eyes and heads enable us to look around without turning our backs on anyone. This, together with all the other kinds of embodied distinctions which are available for investing with interactional significance, is not available to the blockies. The only way for a blockie to 'glance' is by changing its whole bodily orientation. Accordingly, the blockies are considerably constrained in *just how* they can display their attentiveness to others and also in *just how* they can gesture or engage in whole body movements to aid the design of their own turns or to partake in a finely co-ordinated stream of talk. What perhaps is remarkable is that the minimal embodiments offer any interactional affordances at all. Nevertheless, introducing articulations (in the physical sense!) to the embodiments does seem to be worthwhile: e.g. so that a 'looking-around' can be distinguished from a 'turning-away'. Current VR systems essentially treat action in a virtual world as a matter of navigation or object manipulation. However, in CVEs where participants are interacting with one another, perhaps one should consider the direct support of actions of a 'higher-order' than mere movement, actions of *social interactional* significance (like approaches, turnings, glances and maybe some under collaborative control like 'form a circle' and so forth). In future work, we wish to study whether such higher-order actions can be sensibly added to the repertoire of interaction techniques available to the blockies and participants who 'inhabit' them. In this way, we may also be able to help users employ gestures to aid the concurrent construction of their turns - something which is currently problematic.

Of course, adding any further complexity to the blockies has to be reckoned with in the light of technical issues such as computational and network-transmission performance. We feel, though, that a viable and systematic research strategy for developing useful CVEs is to incrementally add further sophistication to very simple embodiments *as and when* analysis reveals that it is called for in the support of social interaction. Interestingly, this goes against the grain of many VR research trajectories which are devoted towards photorealistic body renderings and whole body movement detection. But, unless the social interactional significance of the body is understood, such developments may be not only unduly computationally expensive (especially when one considers distributed collaborative VR systems) but also lacking in social scientific motivation.

References

1. Benford, S., Bowers, J., Fahlén, L., Greenhalgh, C. & Snowdon, D. User Embodiment in Collaborative Virtual Environments, in *Proc. CHI'95*, ACM Press, 1995.

2. Benford, S., Bowers, J., Fahlén, L., Mariani, J. & Rodden, T. Supporting Co-operative Work in Virtual Environments. *The Computer Journal*, 37, 8 (1994), 653-668.

3. Bowers, J. & Churcher, J. Local and global structuring of computer mediated communication. In *Proc. CSCW '88*. ACM Press, 1988.

4. Bowers, J. & Pycock, J. Talking through design. In *Proc. CHI'94*, ACM Press, 1994.

5. Gaver, W. The Affordances of Media Spaces for Collaboration, in *Proc. CSCW'92*, ACM Press, 1992.

6. Goffman, E. *Forms of Talk.* Blackwell, Oxford, 1981.

7. Goodwin, C. *Conversational Organization.* Academic Press, New York, 1981.

8. Greatbatch, D., Luff, P., Heath, C. & Campion, P. Interpersonal communication and human computer interaction, *Interacting with Computers*, 5 (1993), 193-216.

9. Greenhalgh, C. and Benford, S. MASSIVE: A Virtual Reality System for Tele-conferencing, to appear in *ACM Transactions on Computer Human Interfaces.*

10. Heath, C. & Luff, P. Collaboration and control. *CSCW*, 1 (1992), 69-94.

11. Heath, C. *Speech and body movement in medical interaction.* CUP, Cambridge, 1985.

12. Heath, C., Luff, P. & Sellen, A. Reconsidering the virtual workplace, in *Proc. ECSCW'95*, Kluwer, 1995.

13. Presence. Special Issue on Simulator Sickness, *Presence,* 1.3 (1992).

14. Sacks, H., Schegloff, E. & Jefferson, G. A simplest systematics for the organization of turn-taking in conversation. *Language*, 50 (1974), 696-735.

15. Slater, M. & Usoh, M. Representations Systems, Perceptual Position and Presence in Immersive Virtual Environments. *Presence*, 2.3 (1994), 221-234.

16. Tang, J. & Issacs, E. Why do users like video? *CSCW*, 1 (1993), 163-196.

Practically Accomplishing Immersion:

Cooperation in and for Virtual Environments

John Bowers, Jon O'Brien and James Pycock

Abstract

Collaborative virtual environments (CVEs) employ virtual reality technology to support cooperative work. Building on ethnographic and interaction analyses of CVEs in use, we argue that many and varied activities are required to set up, maintain and troubleshoot CVEs. These activities cross-over between virtual worlds and the real, physical environments which meeting participants inhabit. Thus, an understanding of CVEs must attend to the relations between cooperation within a CVE and for it to be established as an arena for intelligible social action. These findings suggest a social scientifically informed respecification of what it is to be 'immersed' in a CVE.

Introduction

This paper presents analyses of the work done in setting up, maintaining, troubleshooting and interacting within Collaborative Virtual Environments (CVEs). CVEs, where multiple individuals interact with each other in a computational environment rendered by Virtual Reality (VR) technology, have been proposed by a number of researchers for the support of cooperative work. Indeed, there exists in the CSCW literature several descriptions of virtual conferencing [11] and collaborative information visualization and retrieval systems [2], as well as general discussions of issues such as how people are to be depicted (or 'embodied') for each other in shared virtual environments [3], and of the relevance of social scientific studies of work and interaction to the design of facilities within a CVE [4]. It is not uncommon for the proponents of CVEs to claim [e.g. 4] that such environments may support cooperative work in ways which go beyond what is possible using more familiar teleconferencing or 'mediaspace' [8] technologies. For example, much CSCW research has indicated the importance of co-participants maintaining a degree of mutual awareness of each other's activity as essential to the coordination of interaction in the workplace and the smooth running of cooperative activities [5, 6, 12, 13]. Advocates of CVEs within CSCW have suggested that enabling users to be embodied in a virtual environment may support people in deploying their everyday interactional competencies to manage turn-taking without artificial technical means such as 'floor control' policies, and to sustain a sense of mutual awareness of work activity within a shared space or 'common frame of reference', features held to be lacking in unsophisticated video conferencing set ups [4].

Our work aims to inform the development of CVEs through empirical study both of settings where CVEs may be inserted [15] and of the use of prototype CVEs [7]. In this way, we hope that our studies can help sensitise developers to the nature of both real world work and the forms of interaction emerging as native to virtual environments. We also intend to add to the stock of studies devoted to the examination of CSCW technologies in use [e.g. 5, 14]. Ultimately, when both this research and the technologies we study have further developed in their maturity, we hope to be able to assess the claims of CVE research in comparison with other orientations to the support of cooperative work. This paper and an earlier one [7] are a start in this research endeavour.

In [7] we presented some analyses of social interaction within a virtual meeting conducted using the CVE MASSIVE [11]. This meeting involved nine participants distributed over

five geographical sites and three countries with each participant being embodied as a simple geometrical shape (a so-called 'blockie', a simple box polygon with the person's name shown above it, a 'face' on one side, a 'mouth' which opened when the person's audio amplitude exceeded a threshold, and two 'ears' which could be displayed to support minimal gesturing). We collected video tapes of the meeting from a number of the sites and one of us was a participant-observer in the meeting itself. Our analyses worked with detailed transcriptions made of the talk and body movements as recorded in the video materials and revealed some systematic problems with turn taking and participation in such environments. Strikingly though, we also found some familiar coordinations of body movement (e.g. exchanges of 'glances') between participants even though the embodiments may seem very minimal.

To our knowledge, this earlier work is the first to utilise interaction analysis techniques derived from the Conversation Analysis (CA) tradition [e.g. 16] to study virtual environments. However, as we mostly restricted ourselves to the analysis of recordings of audio/video output from the machines running the MASSIVE client-visualisers, there are aspects of CVEs as a technology for CSCW which we were not able to speak to. While we were able to discuss many features of the action and interaction within the *CVE itself*, we were not so readily able to show how events and activity within the CVE relate to the 'real world' activity *outside the CVE* and in the (physical, not virtual) settings where (humanly, not computationally) embodied users were working with the system as well as managing the whole host of other contingent activities which constitute their everyday working lives (e.g. acknowledging others who might pass through or answering the 'phone). In short, we were not able to capture fully how action and interaction in the two worlds (real and virtual) might inter-relate from an analysis of video materials of the CVE alone.

Accordingly, the current paper draws more on ethnographic materials (presented through analyses of field notes particularly concerning activity outside the CVE) but supplemented, as required, by interaction analyses (based on transcriptions of activity within the CVE). We believe that such a combination of methods is appropriate to address research questions surrounding the relationship between real and virtual worlds when considering CVEs. Thus, our approach here in combining methods is in line with Anderson, Hughes and Sharrock [1, p.136] when they write: "Methods ought to be chosen in the light of the problem rather than vice-versa. The methods available to any discipline are not canonical forms of inquiry but only stocks of methods in hand and there is no reason why their number should not be increased as problems dictate."

Within CSCW, one of us [5] has emphasised before how important it is to study 'the work to make it work'. By this is denoted all the multiple mundane tasks that are necessary to keep technology working, to 'oil the wheels', to keep things 'up and running', to manage all the unexpected contingencies that arise when good research ideas collide with workplace implementation. Bowers' [5] study of the introduction of a CSCW network in a UK government organization indicated that if the work to make it work was thought to be excessive, this could be reason for abandoning innovative technology even if individual users were well satisfied with it and even if benefits were being experienced. In terms of our current studies of CVEs, it is important to study how work and interaction in real *and* virtual worlds inter-relate for otherwise 'the work to make it work' may be systematically and inadvertently ignored. Noting 'the work to make CVEs work' is important at this stage of CVE research—we believe—because it is very easy to be swept away by the excitement of VR and forget the extensive work required to make such things as virtual meetings happen at all. As CVEs begin to move outside of research labs, it behoves us to document the methods employed by virtual meeting participants, by developers themselves, indeed all concerned, to manage the contingencies involved in actually making CVE technology work. To document *those* methods *our* research methods need to be sensitively reconfigured.

In what follows, we shall document some of the methods used in setting up virtual meetings, in maintaining them, in troubleshooting problems as they emerge, and in artfully managing the relations between activity in real and virtual worlds. In the conclusion of the paper, we examine the implications of our work for the design of CVEs and suggest—in line with our empirical studies—a redefinition of a central concept much used in VR research: immersion.

Setting up the meeting

The virtual meetings in CVEs that we have observed are far from spontaneous affairs in that much planning and preparation occurs in advance of the connections actually being made and the virtual environment created as an arena for the meeting. The dates and times of meetings[1] are planned by means of *other* technical arrangements (the research group's Email list), and on many occasions an agenda is drawn up in advance. Whilst procedures seen as necessary for the successful completion of a task such as 'holding a meeting' might be thought of as straightforward—simply turn up, listen and contribute where necessary—our observations of meetings held in CVEs lead us to contend otherwise drawing attention to the skilled judgement and effective co-operation displayed by those observed in the set up of the meeting. Indeed a great deal of collaboration between individuals in the real world is required in order to collaborate successfully in the virtual world.

Immediately obvious from the following field notes is the exercising of judgement and skill by PC in making the technology do the work asked of it.

```
Close to agreed meeting time (agreed via project's Email mailing list)
PC is checking network activity of machines that will participate in
the meeting. Notes activity and thus starts up MASSIVE.
Problems with audio... there are other embodiments present but no au-
dio: dead silence
PC using text channel... asks for help.. no text response from any one
else
'tempted to restart it, but I don't want to lose this... I think the
audio's independent'
Attempts to fix audio - no joy
```

PC shows a considerable skill and resourcefulness in the face of contingencies that emerge to make setting the meeting up problematic: he checks the network activity of other machines he knows will be used in taking part in the meeting, recognises activity consistent with running the MASSIVE server and proceeds to launch the client on his own machine. It is interesting that PC begins by using machines' activity as a basis for inferring the actions of others and coordinating his responses. This form of monitoring is non-intrusive for others and compatible with an intention to communicate primarily via the virtual environment once established.

We see that the first problem PC encounters is that his audio channel is not operating successfully and that he reacts by asking for help via MASSIVE's text interface, a facility which presents a 2D plan view of the CVE using ASCII characters and allows the exchange of text messages. This request for help is entirely consistent with the common strategy to 'repair' problematic situations with a given element of the CVE through the

[1]It is worth noting that the dates of these meetings are drawn up in order to accommodate the everyday real world activities of participants to the meeting as we see later in a later example where BB asks if a meeting can occur at a certain time because lab assessors are coming around and an ongoing VR meeting will provide appropriate evidence of the success of their current research. The use of this technology is, indeed, resolutely embedded in the day to day concerns of participants.

use of other elements of the system which are more reliable at the moment (in this example the audio and text interfaces respectively). We will also see that other users are aware that such a strategy is commonly pursued and thus monitor the full range of interfaces when problems occur.

In this instance, as no-one responds to him he considers his course of action, and subsequently infers that the presence of others is not sufficiently stable to assist him, and that his faltering audio probably requires a local solution at this point. He therefore attempts to fix the audio first by altering various of the CVE's audio settings and then by fiddling with the audio jacks in their sockets.

Such activity clearly relies upon PC's ability to make judgements about the activities and knowledge of others (their lack of response to him is not interpreted as rudeness or unwillingness to help, but rather as signifying continuing technical problems with both server and client) and relatedly his skilled ability to draw conclusions about the nature of the technology he is using (he knows, for example, that he need not necessarily quit MASSIVE's visualiser in the first instance when he is having audio problems, the two are relatively independent and juxtaposable elements of the system). Establishing the meeting is by no means a simple activity: local knowledge and tacit skills have to be employed in the process.

In addition to the skill and judgement of single users in isolation, *before* groups of users collaborate in holding the virtual meeting, a considerable degree of co-ordination, both between separate users at the same site, and between users at sites remote from one another, is required to set up the meeting in the first place as we can see from the following excerpt from field notes:

```
BH asks if everyone's there who should be?
PC - JF's gone to loo!
BH starts meeting... PC focuses on him, initially
BH gives progress report - asks PC for Lancaster report
JF enters
PC says to the meeting that he will shortly pass over to JF, "who is
just starting up his own body": BH passes on to RK (both other parti-
cipants in the VR meeting)
JF starts up 'which world are we in?' PC tells him - appears to 'kill'
PC who loses audio.
PC turns to JF: 'you've disappeared JF, can you hear anything?'
JF - not yet
PC - tries audio: then 'are you using the old version, JF?'
JF - Don't know... what's speech supposed to be on?'
PC - G
JF - Right
PC - have you got sound?
JF - yeah, got it now
PC - you've probably killed mine: turns to machine, pauses saying
'hmmm' and flicks between windows -  notices new text in text window
reads text: 'SN [at Notts] says leave the room and come back in again'
PC types in text window: 'OK back in min'
```

They discuss possible solutions to the audio problem, and to the likelihood of JF's launching of MASSIVE having 'killed' PC's embodiment. Eventually a participant at another site offers them advice via the text channel as to how their difficulties might be overcome. Such interactions, not at all untypical of setting up a meeting in a CVE, make clear the co-ordination required of participants' activities as they undertake the set up. This co-ordination is based around the routine deployment of sets of skills and local

knowledge (who knows what, who's busy, who's worth asking about 'x') and is of course entirely characteristic of much collaborative work. What is worth noting here is the extent to which this collaboration and co-ordination is required *for the purposes of establishing* a collaborative virtual environment.

Thus the very existence of the VR world depends on the local and sometimes distributed management of real world resources and infrastructures. This form of social organisation is also required in order that the meeting, once established can be maintained. Participants in the meeting make use of the resources to hand in order to order the real and virtual worlds, and it is the work required to achieve this orderliness which is the subject of our attention in the next few sections.

Maintaining the Meeting

The meeting in VR cannot begin until 'everyone is present' but when this has been achieved much remains to be accomplished to order the conduct of the meeting. Our interest is in how meeting participants collaboratively organise the meeting and practically achieve and maintain its orderliness as best they can.

Forming a Group

There are occasions where participants are able to make use of the facilities MASSIVE offers to embodiments in the virtual world. In the following example the task in hand is apparently a more straightforward one than holding a meeting—it is simply to gather in a circle to prepare for a group activity:

CT suggests gathering in a circle. Some complaints... notoriously difficult to manage. Problems of orienting one's own embodiments to those of all others. Some attempts at first, people's embodiments run through each other.

Someone from Site B explains that you can switch to a bird's eye view and see your own and other embodiments from directly above, much easier to orient embodiments in relation to others.

Initially users struggled to co-ordinate their position in relation to other users, before it was pointed out that they were able to 'leave' their embodiments and look down on them from above, making such co-ordination much more manageable. Here the ability to exploit the 'freedom' of viewpoint in MASSIVE contributed to the successful completion of a task by making all participants aware of the work of others. The example also indicates that the form which embodiment in a CVE should take (whether one 'looks out of' the blockie's eye or prefers an 'out of the body experience') is a contingent and activity specific affair.

Trustable Resources

Social actors in the situations we have observed make use of the resources at hand for action and interaction in order to establish and maintain the virtual meeting. Clearly it is intended within the CVE that users' embodiments are one such resource—a means for those participating in the meeting to make their actions publicly available—so that others can draw interactionally significant inferences from them and in turn use their own embodiments to respond accordingly.

Observation of use of the MASSIVE CVE makes plain that situations frequently arise when user embodiments make just such inferential activity problematic, since activities in the two worlds are not 'in alignment'. By this we mean that inferences drawn from the particular state or orientation of an embodiment within the virtual world are not adequate to the actual state of the embodied user and thus do not aid in accomplishing the meeting's business—in the following example its closure.

JF & PC observed at Site A
BH at Site B
End of reports: BH asks if there's anymore DEVRL stuff that anyone wants to talk about... silence... [no way for people to shake heads]... just have to wait for response or lack of it!
PC looks over shoulder, checks out JF's machine
BH attempts to wrap up: 'have we had enough?': 'MASSIVE saturation?'
More silence in meeting
JF asks question - when is next meeting? silence no response. No way of knowing if people have heard - People start playing around with embodiments: flying, sinking through floor, 'doing Jaws impressions'
After 5 - 10 minutes
PC asks if that is it... everyone continues to wander around “someone has said is that it? and now everyone's just wandering around”
More silence in the meeting

Here we see the problems associated with the participants' limited ability to use the embodiments to make others aware of their work and furthermore to infer from others' embodiments their activities. Unlike a meeting in the real world, participants are not able to make plain certain highly meaningful nuances of their perspectives on the virtual meeting—in this case the silences, lengthy pauses between comments and unanswered questions present real problems to participants in expressing their own and inferring others concurrence (or lack of it) with the proposed closure of the meeting ('have we had enough?': 'MASSIVE saturation?') [see also 7]. As PC said after one meeting:

“It's really difficult talking to people on that thing because you'll say something, and you've no idea whether they've heard you, and then you don't know whether they've heard you and they're just ignoring you, and then they don't know what to say, and they heard you and they've heard something which you can't hear and so you haven't replied. You've no idea...”

Precisely what PC identifies as missing, then, are elements to do with supporting the awareness of work amongst collaborating individuals; at issue, then, is the variable ability of the embodiments to articulate successfully the activity of those participating in the meeting, for the practical purposes of coordinating action and interaction. Here we wish to relate such issues back to the idea above of user embodiments *as a resource* for such coordination and introduce the notion of the 'trustability' of these resources in CVEs. The idea of a 'trustable' resource speaks to members' ability to rely upon, for all practical purposes, a given element of the system as a coherent, consistent and publicly available resource for the coordination of action and interaction.

Corpses and Other Untrustworthy Bodies

As we have seen, on such a definition user embodiments are not always 'trustable' resources, and as such can create problems with regard to the orderliness of activities within the meeting. Consider the transcript below which shows AB interrogating a 'corpse' (the term used by the researchers to denote an embodiment which endures in the virtual world even though, for example, the client program run by the embodiment's user has crashed).

AB: can you hear me kai? (2.5) hello?
CG: he may be a corpse now. he's disappeared off my text (as far as)

'Corpsing it' is the most perspicuous form that 'untrustworthy bodies' can take. However, it is not the only form. The next transcript shows SB interrogating a body, or possibly two, who are unable to 'hear'.

```
SB:   are we going to have another one of these?
AB:   we ought to have one with dive.
SB:   yer (0.5) what's-what's-what's the status of that stuff? (1.0)
      he says looking at lennart and kai.
(5.0)
AB:   ah:: lennart can you hear us?
(5.0)
SB:   (hi) if you look at the text window (.) kai says he's not hea-
      ring anything
```

Here, within the meeting, SB and AB discuss the possibility of having a follow-up meeting using the DIVE system, prompting SB to ask about its status. This remark can be heard as directed to LF and KJ as they are the researchers present in the current meeting who are concerned with this other system. However, LF and KJ offer no contribution about such a future meeting when DIVE is first mentioned nor do they when prompted by SB. Interestingly, SB—during the one second pause after explicitly asking about the status of DIVE—moves his blockie directly up in front of LF and KJ. The video taken from KJ's machine clearly shows SB looming large in front of KJ and SB's 'mouth' moving. Shortly thereafter—and while AB is following up to question about LF's abilities to hear—KJ starts typing in a message to the text window, which SB then notices and reports on. This example indicates a number of points of interest. SB uses the embodiment to aid—indeed reinforce—his selection of LF and KJ as the addressees of his question about DIVE. That he has done so and that he has spoken is noticed by KJ and prompts his use of the text window to explain—in advance of his non-response being interpreted in any other way—that he simply cannot hear. That is, KJ juxtaposes communications media to troubleshoot the failing audio and to ensure his conduct (or absence of it) is not interpreted as indicating rudeness[2] or absence from his terminal. In short, when the embodiment fails as a trustable resource for interaction, other media and methods—anticipated as being more trustable at the moment—are resorted to, if necessary, outside of the graphical CVE.

Through such strategies the 'maintenance' of the orderliness of the meeting is something which is constantly being oriented to by users of the system, in part through their displays of attentiveness to other users articulated through their virtual embodiments within the CVE and also through their utilisation of other interfaces to the MASSIVE system in order to contextualise their understandings of the activities (or lack of them) of others and thus draw inferences from embodiments within the virtual world.

Managing Multiple Worlds

Such observations have made clear to us that considerable skill and effort is involved in setting up the virtual meeting in the first place and subsequently maintaining it in the face of variably trustable resources for the coordination of activity. We have been keen to understand the activities of participants both within the virtual environment and within their real world. However, there is a mutually supportive nature to the orderliness of the two. That is, the establishment of a cooperative virtual environment by these participants has produced for each of them an additional social arena beyond their real locale, an arena

[2]Indeed, we argue in [7] that the non-appearance of expected turns at talk (e.g. absent answers to questions) is almost invariably first interpreted as indicating technical problems rather than accountable social behaviour. See also our analyses following the first field note excerpt.

which, like any other is socially organised. As we have seen this social organisation relies heavily on the artful maintenance of the meeting as a technological 'set up' by parties to the meeting. It therefore finds itself inescapably embedded within the everyday contingencies of the socio-technical order of the *real* world of each and every participant and thus we have observed a number of interesting phenomena resulting from the relationship between the participants' interaction in these two worlds.

Exploiting Real Co-location

On a number of occasions we observed participants who enjoyed real world co-location exploit this fact to aid their participation in the virtual meeting. In the following examples, JF and PC are running MASSIVE on different machines in the same room:

```
JF starts up 'which world are we in?' PC tells him
```

In this case JF launches MASSIVE on his own machine in order to join the virtual meeting which has already begun and in which PC is presently participating. JF benefits from PCs combined involvement in the virtual meeting and simultaneous co-location in the real world. In particular JF consults PC about the virtual location of the meeting because the MASSIVE system supports the concurrent display of multiple virtual rooms and other spaces. Later on JF also benefits from PC's co-location in troubleshooting and in monitoring events in order to prevent the need for such troubleshooting work—work which is often difficult given the multi-participant distributed nature of this state-of-the-art system [7].

```
JF covers mic and turns to PC I've got a 'fork fail' on Orchid - they
discuss this in room momentarily and don't pay attention so much to
meeting
They go on to discuss version of DIVE 3 that someone in meeting is
trying to ftp - have they got it? what version is it... ' PC; “it's
the binaries?”
JF asks PC to tell the others he is going to have to restart Orchid
```

Here JF notices a technical problem emerging with his client machine running the CVE. He is able to 'side-step' out of the virtual world and, in the real world, discuss his troubles with PC. Their discussion continues 'out of meeting' but moves around to a topic mentioned earlier in the meeting by another meeting participant (the availability of a copy of DIVE 3). During this discussion JF has been continuing to monitor his machine 'Orchid' and has concluded that he will have to restart his system. JF requests that PC, next to him in the real world, informs others in the virtual world that this is what JF will be doing and that, whatever the embodiment of JF in the virtual world may appear to convey, in fact JF will be temporarily unavailable to other meeting participants. We see then clear benefit for JF in being able to briefly switch attention out of the virtual world and to pressing matters of a local flavour. However we also see that when it becomes clear to JF that he will have to divert his activity for sometime in order to restart his machine he is particularly concerned about possible misinterpretations of his embodiment by those within the virtual meeting. For this reason he—again exploiting his local and real co-location—requests that PC announce his shutdown to the members of the virtual meeting as this event would not otherwise be made plain to them.

Interruptions from the Real World

Indeed we have observed a number of occasions where precisely this type of confusion occurs, that is, situations where other users erroneously infer the availability of a participant for interaction on the basis of their embodiments. In the following example JF, who is participating in a meeting in MASSIVE, is interrupted by a more pressing concern (his

boss) in the 'real world', removes his headphones and turns away from his screen to talk with the interrupter.

```
JF is talking to BB in real world (phones off): PC hears a question -
thinks its for JF... tells him, JF still embodied in virtual world:
has to switch back... hasn't heard question... can't pick up what's
going on
```

Unbeknown to JF a question is asked of him in the virtual world, and it is left to his co-located co-participant PC to prompt him to 'return' to the virtual world. Though turning his attention back to the virtual meeting swiftly JF is nevertheless unable to 'pick up the thread', and as a result he cannot contribute effectively for some time. Thus, the co-location of JF and BB distracts JF from his business in the virtual meeting and as this distraction is not available to the members of the virtual meeting it leads to confusion and a temporary breakdown in co-ordination. Discussing this with PC afterwards, JF highlighted the difficulties presented by presence in both the real and virtual worlds:

```
Yeah like then... er people were talking then and BB was talking [in
real world], so I was talking to BB... you've no idea, you've got two
completely different conversations going on at the same time... you've
got no idea”
```

However, our aim is not simply to show *that* the virtual environment is one among several being concurrently managed but to show *how* such 'multi-tasking' is practically managed individually and cooperatively in this case. We have observed a number of simple techniques deployed by those involved.

Switch Management

We have already observed how JF requests PC to announce his unavailability in the virtual meeting while he restarts his client system. A series of similar phenomena include informal roll calls, reporting on and speaking for absent others, and passing participation to someone else co-located in the real world but who does not have their own distinct embodiment in the CVE. PC refers to this as 'passing over' in the following example.

```
JF & PC observed at Site A
BH & SN at Site B
JF goes off to loo
BH asks if everyone's there who should be?
PC says JF's gone to loo!
BH starts meeting... PC focuses on him, initially
BH gives progress report - asks PC for Lancaster report
JF enters
PC says he will pass over to JF
```

Such phenomena constitute socially organised strategies for maintaining the orderliness of the virtual world in the light of contingencies arising within the real world, skilfully managing switches between the two. We have a number of further examples of occasions when switching between real and virtual worlds is done explicitly so as to provide participants with an accurate sense for what might otherwise appear confusing or 'out of order'.

```
BB enters room (real one!), asks who's there, asks when next planned
meeting is
JF (using MASSIVE) relays question to VR meeting: "BB's here, wants to
know when the next meeting is"
relays answer:
```

```
BB asks are they busy tomorrow at 5?
JF relays question... gets response: to BB - “do you want a meeting?”
BB - no, not particularly... but if there's someone around
JF - (to meeting) is anyone free at 5 tomorrow?
BB wants to demo this to assessors
JF - AF will be there
BB - great.
```

In this case JF informs the virtual meeting that he has, in the real world, BB with him and that BB has a question he would like relaying to those in the virtual meeting. JF then unproblematically switches between one environment and another. Mechanisms to indicate switching between worlds can variably intrude upon ongoing talk in the meeting—for example having to *say* one is just answering the phone or that so and so is leaving may have to be inserted as interruptions in the stream of talk while 'going to sleep' by making one's embodiment recline is used conventionally by the people we have studied to indicate a temporary unavailability for interaction and may be no less noticeable but need not interrupt ongoing talk.

Now compare the next transcript where one participant's anomalous unannounced switch is repaired by others co-located at his site.

```
AB:   <mouth click> .hhh (2.2) <mouth click> .hhh-woahruh i'll go next
      (.) then if no one else is speaking. (0.6) uh i'm adrian bullock
      also from the university of nottingham.
(4.8)
CG:   bizarre-
(4.0)
SB:   he is.
(1.2)
AB:   most definitely.
(0.8)
SB:   anyone else?
```

First, AB introduces himself to the virtual meeting[3]. After this, there is an 'intrusion' from the real world as CG notices anomalous behaviour of the MASSIVE server, which he is maintaining as well as participating in the virtual meeting. His 'bizarre' which is in a sense 'said to himself' but projected to be heard by AB and SB who are in the same physical room as CG is picked up by the microphone and made available to all. Thus, what may have been designed to be overheard by AB and SB to alert them to a potential problem with the local machine which is running the server application is actually projected into the virtual environment as well, where its sense is somewhat obscure. SB and AB regain the orderliness of the virtual world by treating this utterance as if CG was describing AB. Having made CG's remark about the performance of the server in the real world into a joke in the virtual world, SB is then able to reassert the ongoing business of people introducing each other in turn. In this way, a 'switch' from the virtual to the real world which left an inadvertent 'trace' in the virtual one is made manageable and its anomaly repaired by being rendered as a joke at AB's expense (the last speaker to explicitly attend to the business at hand in the virtual world).

[3]This example also demonstrates the sometimes exaggerated use of the audio channel to compensate for embodiments which cannot display a 'readiness-to-initiate-talk' by familiar movements like leaning forwards or raising a hand or whatever [7]. Here AB makes a clicking sound, issues an audible in-breath, pauses for over two seconds, makes an audible out-breath and a vocalization we transcribe as `woahruh` before initiating talk.

Inferring Real Events from the Virtual World

While the real world undoubtedly impacts upon action in the virtual world the relationship is not at all one way. In the following example we see PC notice MASSIVE's performance markedly deteriorate and it is this which prompts PC's actions in the real world.

```
PC comments on the poor performance of the system - someone has
started to back up arb: 'Oh no' - it'll slow everything down'
PC mails him, asks him to delay back up.
PC eventually restarts MASSIVE - informs others via text channel-
'should clear everything out... just make it work'
```

The deterioration in the responsiveness of PC's MASSIVE client leads him to check network activity within his department's computing environment and from this PC realises that someone located elsewhere within his institution has started to back up the local server (called 'arb') which PC is using. This is highly problematic in that it will continue to significantly affect the performance of MASSIVE for this period. To avert such problems PC fires off an Email to the individual his local knowledge tells him is usually associated with the task of backing up 'arb', requesting that the back-up is delayed whilst the meeting is taking place.

The Relation Between Virtual Embodiments and Persons

We have seen how participants manage the relationship between real and virtual worlds to maintain the orderliness of both. The following transcript indicates that the coexistence of activity in virtual and real worlds sometimes requires the artful management of the audio set-up—a matter of importance in its own right. But perhaps more interestingly for current purposes, this leads to a discussion of the contingent and variable relationship that embodiments in the virtual world can have to real world participants and hence to different forms of participation in the meeting.

```
LF:   okay. errm. (2.5) we have err (.) umm (.) a rather a serious
      echo here. an:d erm (1.8) do everyone use headphones and so on?
      i mean it seems (.) some (.) people (.) might (.) be using-umm
      loudspeakers.
(1.5)
KJ:   yes i'm using loudspeakers becos i'm a (higher mind) consisting
      of four people {but i'm trying to keep the microphone covered}
      whenev-whenever i'm not speaking. so (0.5) i shouldn't add (.)
      to the:-uh (.) too much echo.
```

So as not to disrupt the sound in the virtual world, KJ has taken the measure in the real world of covering his microphone as required. KJ has to engage in this elaborate behaviour and not use headphones because there are three other people in the physical room with him looking over his shoulder and observing the virtual meeting without directly participating in it. KJ refers to this arrangement as if he were a 'higher mind consisting of four people' and so his blockie does not stand in the one-blockie-to-one-person-ratio that participants might normally expect. Indeed, inspection of a video taken of KJ and his colleagues reveals that, while KJ has unique control over the keyboard and mouse (and hence of the blockie's movements), all four share in the blockie's point of view and access of received audio. Notably, it is *invariable* in 'round table' self-introduction episodes at the beginning of virtual meetings for participants who are—in some way—sharing their embodiments to specifically note this and inform the rest of those involved. Again, this serves to align expectations about how the relevant embodiments might act and to whom (in terms of real persons) their perception of the virtual world might be made available. It

also alerts meeting participants of the possibility of switching behaviour which might lead otherwise to ambiguities in the blockie's conduct in the virtual world.

Putting It All Together

We close our presentation of empirical materials with a transcript example which remarkably compresses many of the types of phenomena we have discussed in less than 30 seconds of action within a virtual meeting. (This example also uses our conventions for transcribing the movements of the embodiments which we hope will become clear from the analysis of the example. We transcribe on the spot turnings about with ∞---∞ the length and the position of the symbol corresponding to the analogous position in the talk above it. We transcribe translation movements (xyz-displacements) with ^---^. A verbal description of each movement is given just below each movement-transcription. A period (.) at the beginning of a line is used to match up lines of transcribed movement where no body movement occurred with the corresponding line of talk.)

```
SB is facing the blackboard and is closest to it. Others are standing
      looking towards it and SB, who at the beginning of this extract
      has his back to the others.
SB:   and i suggest (.) that (0.2) umm we now do some round table in-
      troductions even though there's no table.  (1.0) i'm steve ben-
      ford from the university of nottingham.
SB:                                  ∞----------------------------
      ------------------------------------------------∞
      .

SB:   turns from the blackboard and faces the others
LF:   .
      .
              ^--------------------------------
LF:   moves towards SB
      ( 7.0 )
LF:   ------^
LF:   continues to move towards SB and passes through him

LF:   (what are you doing?)

():   <laughter>
LF:   ^---------
LF:   backs away from SB

KJ:   lennart har sverigheter att navigera
LF:   ------------------------------------
LF:   continues backing away and comes to
      ( 4.5 )
LF:   ------^
LF:   a halt

LF:   (well-)i was transported by some (      ) person a by- a
      bystander that wasn't very innocent
      (0.6)
SB:   yeah. (0.3) but who are you?-nd where are you from?

LF:   ( 4.3 ) (asking me this now?)
KJ:     ∞----∞
```

KJ: turns to face LF

SB: yeah.

LF: (1.6) umm (2.0) i'm um (.) lennart fahlén from the swedish institute of computer science

Initially in this example, the meeting participants orient their embodiments so that they can see the virtual blackboard upon which the agenda has been posted. SB who is acting as the 'chair' of the meeting suggests that the participants introduce each other in turn while he turns away from the blackboard to face the group. Once he has completed his turning movement, he introduces himself. Interestingly, SB's movements are coordinated with his own talk. He starts the movement just after a short pause and during a hesitation (marked by (0.2) umm) and finishes the movement at the end of his next (a one second) pause. During SB's introduction, LF commences a movement straight towards SB, passing right through him. After making this unexpected movement, LF asks what are you doing? which can be heard as a peculiar remark as it is LF himself who has made the unexpected movement. These events are greeted with laughter and KJ remarks in Swedish that LF is having problems navigating. During the laughter and KJ's remark LF backs away from SB, comes to a halt at approximately where he was before this aberrant movement and explains in humorous terms that his embodiment was moved by someone else in the real world environment (a bystander that wasn't very innocent). Indeed, another person in the lab with LF, MS, was giving him some instruction in how to move his embodiment and had taken control of the mouse to demonstrate forwards movement. This was plain to JB, the other meeting participant present in the same lab along with LF, but would be completely opaque to the virtual meeting participants at distant sites. They merely see LF's embodiment moving unexpectedly in a manner uncoordinated with the ongoing business of the meeting (SB's self-introduction), followed by LF bizarrely asking what are you doing? This strangeness is resolved once we realise that the referent of 'you' is MS, the person in the real-world lab who had been giving LF instruction in moving the embodiment, and *not* any virtual meeting participant. Prior to this moment MS had been involved with the setting up activities in the lab in support of the virtual meeting though he was not going to be a meeting participant. As a 'helper', he had not been attending to the fact that the virtual meeting had begun its business and that his demonstrations of how to use the embodiment could be interpreted as 'movements-in-the-meeting' and not just 'experimental-try-outs-for-LF's-benefit'. A more 'normal order' is only regained once LF backs off to a 'discrete distance' (now controlling his own embodiment) and gives a humorous account of what had happened enabling SB to reassert the meeting's business ('round table introductions'). LF who had been attentive to his problems coping first with navigating his embodiment and then with handling his instruction (and the sense that meeting participants might give it) not surprisingly requires confirmation from SB of what is required (asking me this now? after a long pause of over four seconds). However, SB's initial attempt to elicit a self-introduction from LF (but who are you?-nd where are you from?) is enough for KJ to orient his embodiment to face LF, a movement he completes just before LF requests confirmation.

In this example, although the simple geometrical shapes which compose the blockies may constrain the subtlety of nuance that can be portrayed, we can see recognisable movements of the embodiments of social interactional significance ('turnings to face others', 'turnings to display attentiveness', 'movements to invite participation'). Indeed, they are familiar enough for an *unexpected* movement to be recognised as such and be accountable. An intrusion from the real world, due in part to a non-participant's non-attentiveness to events in the virtual world, had to be repaired for the benefit of those within the virtual

world both by means of a verbal account and a restitution of a normal body position. Once this is achieved order can return to the virtual world.

Conclusions

We have seen how participants' interaction with others in the real world benefits and impacts upon interaction in the virtual world and vice versa. In this way it is evident that while the social arena of the virtual world requires its own organising it must also be practically managed with reference to events in the real world. Participants are seen to display various skills, methods and resources in handling the dynamic and variable location of their activities within these arenas. Each participant is then simultaneously operating in several 'worlds'—some real, some virtual, some local, some nearby, and some distant. Furthermore, while the alignment of these worlds is practically managed during the real-time of the meeting, for things to come together at all, much work is required in assembling the worlds and tending to them in the first place. In short, CVEs may well support cooperative work; they certainly require it.

Practically Producing Trustable Worlds and Embodiments

We have suggested that how 'trustable' a resource is matters to its deployment for interactional purposes in a CVE and, in the introduction, we noted the central importance embodying users has to the philosophy of CVEs as an approach to CSCW. Well, are the embodiments trustworthy as the bearers of meaning and significance?

Yes indeed, they can be. *If* activities in the real world are adequately aligned with the embodiment's activities in the virtual world, *if*—for the purposes at hand—the embodiment adequately displays its attentiveness to ongoing activity, *if* participants are adequately mutually focused on a common thread of activity, and *if* the machines are working, *then* the embodiments can be relied upon both by their 'owners' and by the others as trustworthy resources for social interaction. Under such circumstances even just a subtle turning or an approach can influence the course of social interaction (e.g. by selecting who is next to speak or eliciting further details) or be closely meshed with it (e.g. by timing movements in close coordination with those of others or their speech [7]). However, *if* users engage in activities in the real world which mean their virtual actions are either inappropriate (e.g. if a not-so-innocent bystander gives not-so-innocent ad hoc instruction) or non-existent when expected (e.g. if a call of the 'phone or of nature has had to have been answered), or *if* the embodiment cannot adequately portray a user's attentiveness or readiness to participate (e.g. if an embodiment has 'corpsed'), or *if* activity devolves into multiple threads (e.g. as some users are engaged in troubleshooting local problems, while—at the same time—others complain about the audio, and yet other parties are trying to focus the gathering on introducing themselves), or *if* a machine has crashed, *then* the embodiment cannot be trusted as a user's representative in a virtual world and other methods will have to be resorted to (closely attending to or adapting one's use of the audio link, scrutinising the text window, launching unix talk or making a 'phone call).

Evaluating and Developing CVEs

On this view, the utility of the embodiments is *not* given by technical arrangements alone. Their usefulness as bearers of interactional significance has to be worked up, worked on, maintained and repaired in the light of ongoing activity in both real and virtual worlds. This means that the evaluation of the embodiments and—more generally—of CVEs as an orientation to CSCW cannot be made in any simplistic way. Absolutely not. Naturally, analyses of the sort we have presented can be taken to have implications for the design and redesign of embodiments and virtual worlds for cooperative work. In [7], we detail a number of possible ways in which the design of the 'blockies' might be enhanced to sup-

port social interaction more effectively. However it would be too crude a development strategy to suppose that technical remedies should be matched to every social interactional problem. For example, we have indicated that—at systematic moments—participantsmay have a sometimes exaggerated recourse to the audio channel to compensate for the embodiments' constrained abilities at displaying attentiveness and readiness for interaction. One should not *automatically* assume that this should mean that the embodiments *must* be redesigned accordingly. For many purposes, this shift of activity to the audio channel may be entirely adequate with the presence of the virtual world still yielding an effective composite environment for cooperative work with potential advantages over alternative solutions.

Furthermore, the supporting real world activity that virtual worlds require for them to be set up and maintained may well be justified in the light of potential advantages of virtual meetings and the fact that the systems we are studying are still research systems, rather than turnkey commercial solutions. In short, 'the work to make it work' [5] needs to be calibrated against the anticipated merits of using new technical arrangements in any definitive judgement of the value of different design options and development strategies. At this stage of CVE research, it would be premature and misguided to offer any such definitive evaluation. What we can do now, though, is to indicate the kinds of criteria which might underlie such an evaluation in the light of the phenomena we have observed. For example, 'the work to make it work' needs to be manageable, the relationship between different communications media needs to be such that technical troubleshooting can be supported and social interactional troubles repaired as they arise, the embodiments need to afford displays of attentiveness and enable mutual awareness of activity to be sustained but, if they cannot, then artful juxtaposition of other media should be available to adequately compensate. And so forth.

In this regard, we are especially concerned that the reader does not get the impression that we are negatively evaluating the MASSIVE system which has been used in most of the virtual meetings we have studied. Indeed, the fact that participants in the virtual meetings *can* achieve an orderliness to their activity within the CVE testifies well to its design. That it provides a number of 'interfaces' (graphics, text, audio) which can be juxtaposed to troubleshoot problems and that the embodiments *can* be invested with social interactional significance [see also 7] are notable design achievements. To document, as we have done, 'the work to make it work' and to show that it consists of many and varied activities is *not* to downgrade the technology involved. It would only seem so to those who tend to forget such work, ignore it or wish it away. All technologies require set-up, maintenance and repair work. It is idealistic to believe otherwise[4]. What matters is that such work is possible, manageable and commensurate with anticipated benefits. For the researchers we have studied, this is undoubtedly the case. It is an open question whether that would be true in non-research settings and this is why we complement studies such

[4]For this and related reasons, we do not think that it is correct to argue that our analyses are of little general interest because of the research-oriented, non-commercial and in-development nature of the systems we have studied (as one reviewer claimed). First, the systems we have studied *are* state of the arts in CVE development, so if the use of CVEs is to be studied empirically at all right now, it will involve the ones we have examined or ones like them. Second, as we have emphasised, the 'work to make it work' doesn't go away with more 'mature' systems, though it may be overlooked or we may become used to just accepting it or leaving it for 'system managers'. [5] shows that considerable, and typically unrecognised, supportive work is required for commercially released CSCW and groupware systems. To dismiss such work as mere 'system management' would be superficial and, in the case of CVEs, miss critical issues to do with the *relationship* between activity in real and virtual worlds—issues which may ultimately be critical for the user-acceptance of commercial CVEs.

as that reported in this paper with field study in settings where candidate end-users might work [15].

As a final warning against premature and crude assessments of CVEs for CSCW (or of taking our paper as offering such), let us emphasise also that such criteria need to be relativised against the demands of specific applications of CVEs and the specific workplaces which they are intended to fit into or transform. Our studies have been of CVEs in early use in research communities to support research meetings. Accordingly, our orientation in studying CVEs has been to investigate the work both inside and outside the CVE which is necessary to accomplish 'having a meeting'. We have *not* presented studies of other possible applications of VR within CSCW (e.g. cooperative information visualization and retrieval [2]), nor studies of usage within non-research environments. While our studies may sensitise future research into the usage of CVEs in such settings by, for example, suggesting the phenomena which *may* occur and be significant, this is no guarantee that they will occur in these other settings and be significant for them. Only further study can reveal that.

Virtual Reality for CSCW Reconsidered

At the outset of this paper, we noted some of the claims made by proponents of VR technology for CSCW and the potential advantages CVEs may have over, say, crude approaches to teleconferencing. We find ourselves in complete agreement that supporting participants' mutual awareness of each other's activitiesand—relatedly—enablingparticipants to display their activities and attentiveness to what's going on in subtle and differentiated ways are essential to the potential of CVEs and indeed to much CSCW technical development. What our work points to is the *totality* of activities involved in attaining awareness of others and of displaying one's own activities to them. Some of these are mediated by virtual embodiments and what they afford but many other resources are drawn upon in artful ways by participants to support mutual awareness and an orientation towards the sociality of the ongoing work (in the cases we have studied the ongoing work of 'having a meeting').

Indeed, being able to deploy these 'other resources' (from a knowledge of the local computing environment through to adapting one's speech action over the audio link) in an artful fashion is often in many senses a *pre-requisite* of investing the virtual world and the embodiments with the sense that is required for cooperative work within it and with them. Accordingly, when it is claimed that CVEs can in principle support cooperative work in ways difficult to achieve with alternative technologies, we take this as a claim most appropriately assessed in the light of *all* the work *both within and without* the virtual environment narrowly considered. CVEs should not be criticised (nor prematurely celebrated) on the basis of *only* what can be designed into the virtual environments or occurs within them. This argument, then, suggests that there is an opportunity to rethink the design challenges in CVE research: *one should be designing for two worlds not just one.*

Similarly, methods of empirical investigation must be carefully configured so as to minimise the chances of superficially analysing activities of significance to system use, or missing them altogether. If this involves the adoption of a repertoire of methods (as in our combination of ethnographic and interaction analysis), then so be it.

Reconceiving The Concept of Immersion: Immersion as a Practical Accomplishment

At some juncture in this paper the reader may have considered the objection that much of our research seems to depend upon us having studied 'desktop' VR systems and not 'immersive' ones where, for example, head mounted displays (HMDs) and body position and movement sensors, perhaps combined with projected video images of faces (and so

forth), would be used to both more thoroughly portray the user within a virtual world and give the user a richer experience of being present. Within such arrangements for example, or so it may be argued, users do not have the option of being unaccountably absent from their 'terminals' as they no longer interact with a (real-world) terminal in order to act within a virtual world.

First, of course, it is an open empirical question whether the phenomena we have observed can also be found with such immersive systems. However, we would be most suspicious of claims that the mundane familiar world could simply be ignored in the appraisal of such systems, as if this forgetting could be underwritten by the allegedly more intimate connection between human bodies and their virtual counterparts. Work will still be required to set up connections, maintain them, repair technical troubles and so forth (and with the delicacy of much current immersive technology, maybe more so). And—unless participants lead an entirely immersed life—the physical world is bound to interrupt at some moment or another[5].

Our *second* response to the objection that we have confined ourselves to non-immersive, desktop virtual worlds is to urge a reconsideration of what is meant by 'immersive' in such contexts. We would suggest the utility of understanding 'immersion', *not* in terms of different technical arrangements (e.g. HMD versus screen-based presentation), *nor* in narrowly defined perceptual or psychophysical terms (e.g. is the virtual world all that one can see/hear? to what extent does the virtual world take account of human perceptual systems?), *but* as a practical accomplishment brought off through the work done in giving social activity an orderliness in the virtual world. As such, immersion would need to be understood as involving *relations* between one's real-world activities, one's virtual world activities, one's abilities to display one's activities in either world to others, relations which are ongoingly achieved, maintained and (if necessary) repaired. We would suggest that CVE research in CSCW may benefit from a sense of 'immersion' which recognises the sociality of activity within worlds both real and virtual and not merely the relationship of a single user's perception to a virtual world, or which holds that immersion—in anything other than a minimal sense—is achieved merely by technical arrangements (HMDs, 'cybersuits'). On the view we are developing, *once* actions in the real world are adequately aligned with actions in the virtual world, *once* one's attentiveness to ongoing activity in the virtual world is appropriately reflected in what one does with one's embodiment and what the embodiment affords, *once* local technicians are persuaded to delay their weekly backup, *then* one can say that one's participation is adequately immersed.

Ordering the Virtual World as Mundane

A great deal of excitement surrounds VR and, in some quarters, this has turned into hype and exaggeration. For each published account of a working VR system, we read many more announcements that VR comprises a technological 'brave new world' or that the possibility of 'electronic embodiment' ushers in new forms of human subjectivity. Our research makes a different impression. We see researchers and meeting participants engaged in their day-to-day business, exploiting local knowledge, mobilising everyday skills and competencies, artfully managing contingencies and problems as and when they come up as best they can. We see ordinary interactional competencies (methods for managing turn taking, displaying attentiveness and orienting bodies, using another means if one fails) deployed so as to make less familiar, less trustable arenas for action more recognisable and reliable than they otherwise might be. In all this, the virtual world is but one domain and the management of multiple arenas appears in many ways a normal and

[5]Even the most deeply immersed of contemporary cybernauts need to sleep, feed and the rest, no matter how minimally. Natural functions are not yet computable!

unexceptional task and in that sense mundane—which is not to say that it requires no skill, indeed quite the reverse. But the skills it requires are those of appropriately competent researchers and meeting participants, rather than 'cyborgs' or 'postmodern subjectivities'. Activities—be they in the real or virtual worlds remain *worldly*. How could it be otherwise?

References

1. Anderson, R., Hughes, J. & Sharrock, W. *The Sociology Game.* Longman, New York, 1985.

2. Benford, S., Snowdon, D., Greenhalgh, C., Ingram, R., Knox, I. & Brown, C. VR-VIBE: A Virtual Environment for Co-operative Information Retrieval, in *Proc. Eurographics'95,* Maastricht, The Netherlands, 1995.

3. Benford, S., Bowers, J., Fahlén, L., Greenhalgh, C. & Snowdon, D. User Embodiment in Collaborative Virtual Environments, in *Proc. CHI'95,* ACM Press, 1995.

4. Benford, S., Bowers, J., Fahlén, L., Mariani, J. & Rodden, T. Supporting Co-operative Work in Virtual Environments. *The Computer Journal*, 37, 8 (1994), 653-668.

5. Bowers, J. (1995). Making it work: A field study of a 'CSCW Network'. *Information Society,* 11 (3), 189-207.

6. Bowers, J., Button, G. & Sharrock, W. Workflow from Within and Without, in *Proc. ECSCW95,* Kluwer, 1995.

7. Bowers, J., Pycock, J. & O'Brien, J. Talk and Embodiment in Collaborative Virtual Environments, in *Proc. CHI'96,* ACM Press, 1996.

8. Gaver, W. The Affordances of Media Spaces for Collaboration, in *Proc. CSCW'92,* ACM Press, 1992.

9. Goffman, E. *Forms of Talk.* Blackwell, Oxford, 1981.

10. Goodwin, C. *Conversational Organization.* Academic Press, New York, 1981.

11. Greenhalgh, C. and Benford, S. MASSIVE: A Virtual Reality System for Tele-conferencing, to appear in *ACM Transactions on Computer Human Interfaces.*

12. Heath, C. & Luff, P. Collaboration and control. *CSCW*, 1 (1992), 69-94.

13. Hughes, J., Randall, D. & Shapiro, D. 1992. Faltering from ethnography to design. In *Proceedings of CSCW'92* (J. Turner and R. Kraut, eds.). New York: ACM Press.

14. Orlikowski, W. 1992. Learning from Notes: Organizational issues in groupware implementation. In *Proceedings of CSCW'92* (J. Turner and R. Kraut, eds.). New York: ACM Press.

15. Pycock, J. & Bowers, J. Getting Others To Get It Right. In these proceedings.

16. Sacks, H., Schegloff, E. & Jefferson, G. A simplest systematics for the organization of turn-taking in conversation. *Language,* 50 (1974), 696-735.

Getting Others To Get It Right:

An Ethnography of Design Work in the Fashion Industry

James Pycock and John Bowers

Abstract

This paper reports an ethnographic study of design work in the fashion industry. Contrary to many images of fashion design, in this setting, it is essentially tied to organizational and inter-organizational coordination, and the demands of manufacture and supply chain management. Relatively little design work involves artistic drawing, much requires retrieval from databases, data analysis, information gathering and matters which members themselves call 'technological'. Experiences collaborating with developers and the relevance of advanced 3D design tools and Virtual Reality for CSCW are considered on the basis of these findings and in the light of debates over ethnography in system development.

Introduction: Ethnography in CSCW

Ethnographic work has become highly influential in CSCW. Although there are many different species of ethnography [10, 14], most share a commitment to the descriptive, non-quantitative study of the details of everyday work practices in real world settings as they are played out in real-time, with—where at all possible—an extensive period of contact between the researcher and the researched. Within CSCW, ethnographic and allied research came to prominence principally through the detailed description of circumscribed locales such as control rooms [8, 10] where cooperative work is achieved through the systematic deployment of members' interactional competencies and the artful use of (sometimes relatively simple) technologies. More recently, there has been a concern to move CSCW's ethnographic research 'out of the control room' as [9] puts it and into more varied business and industrial settings including those not so readily characterised by the 'focused interaction' of more constrained work domains.

This has had a number of consequences. First, the *organizational situatedness* of cooperative work has become a palpable research issue. Studies of technology and cooperative work in banking [12], central government [3] and management training centres [13] all show how organizational concerns of one sort or another are oriented towards in the conduct of the work. Pushing these points further, Bowers, Button and Sharrock [4] have argued that *inter-organizational relations* need to be appreciated in the study of cooperative technology. Second, there is a concern to increase the *variety of settings* studied moving not only out of the control room but also out of the office and onto, for example, the shopfloor of the print industry [4]. Such settings prominently show how the coordination of cooperative work is often tied to the specific materials of the job.

Finally, there exist a number of pieces reflecting on how ethnographic work might relate to system development. For example, Hughes et al. [9] identify four different roles for ethnography in system development distinguishing 'concurrent' ethnography from 'quick and dirty' and 'evaluative' ethnographies and those three from the 'reassessment of previous studies'. Additionally, Plowman et al. [11] critically review workplace studies, many of which have an ethnographic orientation, and argue that there are tensions in trying to serve the dual masters of providing accounts faithful to the detail of the setting under study and usable design recommendations.

Rather than enter these debates directly, the current paper seeks to inform them through the presentation of a further field study—one we have conducted under the exigencies of a multi-participant academic/industrial project in the UK. The context of the project itself, together with the many and varied kinds of findings that our research has come to point to, has influenced how we see ethnographic work in relation to development as well as our view on debates such as those Plowman and her colleagues initiate [11]. Our study also aims to add to the stock of knowledge in CSCW concerning real world cooperative work while further emphasising the importance of studying settings which raise issues of inter-organizational coordination and where the specific materials of the job matter to how the work is done.

Design Work and Virtual Reality for CSCW

Our fieldwork has been conducted in the UK fashion industry and concerns fashion design. Our interest is in characterising the real world work of design in enough depth to inform the development of systems. In particular, we are interested in the role Virtual Reality (VR) technology might have in fashion design work.

VR technology is being increasingly explored in CSCW. Several researchers have proposed 'collaborative virtual environments', where individuals interact with each other in an environment rendered by VR technology, as an approach for addressing a number of CSCW research problems. For example, there exist virtual conferencing [7] and collaborative information visualization and retrieval systems [2] in the CSCW literature. Recently, we have become concerned with the use of ethnographic research techniques for supporting the design of such systems and examining the usage of prototypes [5, 6] as well as exploring the relevance of 'VR for CSCW' as a general development orientation.

The occasion for the fieldwork we report in this paper has been our participation in the UK project VirtuOsi [1]. Two 'pilot' activities provide a focus for the technical development and social science research in the project. One of these, the 'fashion pilot', directly motivated the fieldwork we report. This was concerned with the development of VR technologies for the fashion industry principally through two prototypes: a virtual catwalk environment and a virtual 3D design studio integrated with a multimedia database. Some work on these environments—at least in conceptual form—had commenced before our participation and formed a part of the overall research programme of one of our collaborators. Initially, our collaborators were concerned to capitalise on the real-time 3D graphics abilities of VR systems with a particular concern to support the fashion industry and design activities within such settings. For them, the support of the 3D visualization of garments—particularly garments draped over humanoid figures and dynamically modelled as those figures moved—was a critical problem to be tackled in fashion design work and VR technologies hold much promise for addressing it. Our ethnographic work has a dual status with respect to these development agendas. First, we seek to *inform* them by providing 'intelligence' about the nature of design work in a real organizational setting in the fashion industry. Second, we have come to believe that our fieldwork findings can *augment* current agendas by exploring further opportunities for the support of cooperative work beyond those that would be serviced by the catwalk and design studio environments.

The Setting

The Company

The company we have studied, which we shall call 'DNK Fashions', specialises in the growing retail market of home shopping selling mainly fashion items but also some hou-

sehold goods and electrical items. It conducts its business by mail order largely through catalogues, dealing direct with customers. DNK has grown four-fold in ten years and has become one of the largest home shopping companies in the UK to sell direct to customers. DNK was established over a hundred years ago, has 30 different catalogue titles, and has gained significant market shares in a number of niche markets such as specialised footwear and large sized clothing. It is important to note, however, that the company does not manufacture goods itself. All manufacturing for DNK is outsourced. As we shall see, this raises issues for the coordination of design and manufacture and the maintenance of the supply chain—especially when each season the company sells products in 30,000 different size and colour options.

The company's business concerns high volumes of low cost goods in low value multiple item transactions involving many millions of customers. This is indicated by the company's database of 9 million customers, its average customer order value of £80, and its distribution of over 2 million copies of its main catalogue. The headquarters of the company are located in the centre of a large northern city in the UK. Entire floors of this 15 storey building are occupied by hundreds of people answering telephones or opening letters. Just under 2000 people work on this site and are involved in activities ranging from desktop publishing to footwear buying, from import arranging to packaging technology. Clearly, a constant and critical organizational problem at DNK is how to manage this (growing) complexity and multiplicity, and how to manage it in a timely way so that the organization's practices mesh with the seasonal rhythms that fashion cataloguecompanies must observe. Indeed, the issues of 'time and scale' and the relation of these to CSCW technologies is a way of thinking about the core research questions of this paper.

The Department

We have spent much of our time studying the work that takes place in DNK's Quality Control department based mainly at the city centre site. In line with many organization's implementation of business philosophies which emphasise 'quality', 'QC' has emerged as a department of some centrality, needing to interact with many other different departments within the organisation. QC is concerned with (i) devising the specifications for products and packaging which are distributed to suppliers—with particular attention to the utilisation of the company's data on sizing, (ii) testing and assessing suppliers' samples and inspecting random garments from all final deliveries, (iii) analysing why customers return garments and (iv) providing specialist computing facilities and skills for producing garment patterns and textile designs. It also generates in various ways information about the behaviour of both customers and suppliers that is used throughout the organisation.

Let us say a little about the core technologies which support the QC department. Central to much of their work is a data management system built for the fashion industry and widely used by garment retailers. This system consists of a forms based database that is supplied with a variety of predefined 'pages' for recording product data such as size measurements, packaging and labelling requirements, and the like. Of these it is the 'size chart' which is most used and most involved in the work of QC. This chart constitutes the product specification given to suppliers. Within QC there also exists three other important systems. The first of these is a specialised Computer Aided Design (CAD) system for producing 'piece' patterns which are laid onto cloth and cut around at the beginning of garment manufacture. This system has attached to it a large digitizing board for the input of shapes and a plotter for printing full size patterns and layouts of multiple pieces arranged in such a way as to maximise the efficiency of cloth usage. The second system supports scanning in of images and textiles, the modification of these images and the output of results in the form of a colour laser print. A final system of note is a networked database for storing information on customers' reasons for returning goods.

Technologists

Within QC it is 'garment technologists' who undertake the work of specifying particular aspects of garments that suppliers are to manufacture—together with the testing of garment samples (buttons, fabric, seams and the like) and the assessment of quality and safety standards. Technologists at DNK will often have to carefully check, for example, that garments (such as nightwear) contain fire safety labels, that the proportion of different fabrics (wool, cashmere, nylon and the like) specified by the retail company and described on the product labels are in fact matched by the supplier's goods, that washing instructions are appropriate, and the like. The personal expertise of garment technologists is an important factor in their work as they have responsibility for particular ranges of garments and build up detailed knowledge of these products and their suppliers over time.

Buyers

The garment technologists in QC and 'buyers' in buying departments are each responsible for different aspects of a garment's design and production. While the technologists may specify the details of what garments are to be like once they are manufactured, buyers chose the items for the organisation to sell in the first place. The buying activities at DNK are structured into different departments with buyers specialising in certain products, for example shoes or coats, or in certain catalogue titles targeted at different market populations. Buyers in DNK are often responsible for particular lines for 2-3 years before moving on to another product. In this way the buyers also accumulate personal expertise. Buyers are also directly involved in the purchasing of initial stocks of items and placing orders with manufacturers. Indeed, buyers are often key to the relationships between DNK and its suppliers. As one field site member commented to us: "The buyers are gods around here". Considerable liaison between buyers and QC is essential for the smooth management of production as we shall see.

Suppliers

DNK has a pool of suppliers of which approximately 2000 are manufacturing for the company at any one time. These suppliers are distributed worldwide across more than 50 countries but especially concentrated in South East Asia and Eastern Europe. Much of the work of QC involves liaising with suppliers, understanding their problems, persuading them to make the changes wanted by QC, specifying what suppliers should make and even how they should make it. From time to time, DNK organises 'supplier workshops' for some of its large suppliers during which—we were told—"the whole (company) process from QC to Marketing" is demonstrated to suppliers. Such workshops are just the most perspicuous example of what happens every day during production phases—aligning the practices of production with DNK's requirements and deadlines.

The Design and Production Process

Let us give an overview of some of the core activities along the route from buyers' initial selections of designs, through their manufacture, to DNK receiving goods at their warehouse. First, the buyer makes decisions on designs for catalogue garments for future seasons based upon items they have seen being sold perhaps in New York, Milan or Berlin. They may have bought examples of these items on 'shopping trips' and brought them back to the company. Buyers decide changes they wish made to the design of the garment they have bought in order that the design better suits their market population and to avoid copyright issues. For example, a buyer may wish to have a design manufactured in a different fabric, to remove the collar, to add a belt, to change the embroidery and so on. Buyers give their choice of garments from their trips to garment technologists together with instructions for modifications required. The technologists then produce specifi-

cations of that product which detail the measurements of the garment, any important comments and a line drawing of what the modified garment should look like.

Buyers assemble a range of garments to be included in a catalogue for the coming season on the basis of information on customer demand and the like. The specification produced by the garment technologists, together with the number of each size and colour to be ordered and key production deadlines, will be taken by the buyer on trips to see potential manufacturers and negotiate prices. Once the 'buying trip' is complete, legal contracts will be made between the company and suppliers, and the buyer will alert QC to expect samples for testing.

Supplier samples are of two sorts. The first are 'sealing samples' delivered to QC which QC assess and upon which QC authorises a supplier's production to begin. QC checks these samples and will pass or fail them as many times as necessary until one is produced which adequately meets specification. Passed samples are sent to the warehouse for later checking against selections from final deliveries. The second are 'photo-samples' to be worn by models in photographic sessions for the company catalogues. The production of the catalogue is concurrent with suppliers' production of stock with the implication that the company is *committed* to being able to deliver catalogued items from this time. Indeed, it is this feature of specifically catalogue-based retail that is the source of much of the urgency surrounding production deadlines: once the company is committed through the catalogue to supply a garment, there is little scope for turning back.

Knowing Customers and Suppliers

As is commonly the case in the retail trade, DNK have a number of ways of assembling 'intelligence' about their suppliers and customers. For example, when testing the quality of samples received, the number of fails and resubmissions will be recorded per supplier and per garment. In this way, DNK is able to build up knowledge about suppliers in terms of their reliability and timeliness of their deliveries, the quality of their work, and their suitability for particular garments. This knowledge is made organizationally useful in a number of ways. For example, a report on current suppliers is produced when the production stock has arrived (i.e. at the beginning of the catalogue season itself) and is used by buyers when deciding the most appropriate supplier to lodge an order with. Suppliers who are most and least able to deliver samples in a timely fashion are especially noted.

DNK also have a number of ways of finding out about their customers and their buying habits (e.g. customer profiling by recording each transaction) but of all these means, it is perhaps the way in which returned goods are analysed which is of most importance and current interest within the company. Returned goods—and the reasons customers give for returns—are not only informative of customer preferences, they also speak to supplier quality, individual product performance, and are also of interest to buyers.

Returning Goods

It is a feature of mail order retail that customers find that their orders turn out not as expected and they decide to return them. DNK require customers to include a 'returns notice' with details of their customer account number, the items they originally ordered and so on with any returned goods. Importantly, the notice also contains a section (a space for customer written free-text) where customers are invited to record their reasons for returning goods, reasons such as garments being too large, not the colour that was expected, damaged in some way and the like. QC personnel have suggested to us that free-text responses enable more information to be given by the customer and hence become known by the company. In QC, these returns notices are classified into 99 different re-

turns reasons by two quality analysts who aim to process 120,000 such notices between them every six months.

The analysts produce breakdowns of the returns data for each product and each supplier. In this way the company assembles information on, for example, whether a particular supplier is producing a high level of damaged goods, whether a particular item is oversized and so on. It is hoped by DNK that returns data will be useful for making future decisions and even correcting problems during a particular season. In particularly problematic cases, buyers may show suppliers their returns profile on subsequent buying trips, for example, to help convince them that there is a real problem with their goods and their production practices need to be changed.

Standardly a list of the highest returning items for each garment technologist's area of work is produced after the seventh week of the catalogue being available. For each product in this list a breakdown of returns reasons is provided (e.g. 18% too long, 4% not as expected, 63% too short etc) and the technologist will make investigations into these items including perhaps requesting examples from stock, testing fabric weights in the QC laboratory or trying the garments on a model.

The Problem of Fit

One of the early results of returns analysis was to identify a general problem across catalogue items, across suppliers, across buyer's choices, and so on with *sizes*. Easily the most common reason for customers returning goods is that they did not fit properly. As DNK's garments were made to standard UK sizes the company decided to investigate the validity of these size standards themselves. Accordingly, in the early 1990s, DNK conducted its own survey of 710 of (first) its women customers using over 70 key measurements. They found significant differences between the shapes and dimensions of their sample compared to the old 1950s research which the standards were based on, particularly in terms of waist, crotch length and height. The survey also confirmed returns analysis data which suggested that sleeve lengths were too long by revealing, as a manager explained to us, "assumptions that as people got bigger their bones got bigger". In fact many of the company's customers had been smaller sizes earlier in life but had become a larger (clothing) size but, of course, without any accompanying change in the lengths of their arms!

Making Knowledge Work

We have emphasised the importance that knowing customers and suppliers has for DNK in addressing problems of supply chain management, design choice, demand anticipation amongst other issues. However, it is not enough to merely capture information about customers and suppliers, it has to be *made organizationally useful.* Data have to be classified, sorted, analysed and presented to relevant personnel in ways which can inform, recommend or (on occasion) mandate courses of action. For most personnel at DNK, it is not necessarily the task of collecting data that is problematic (though it is acknowledged that size surveys are expensive and time consuming), it is encouraging people to act on the basis of the data which can be troublesome. Making data 'actionable' is a key organizational concern.

We have already seen some of the ways in which returns data is made organizationally useful in the form of supplier reports which can be consulted by buyers and so forth. However, making the size survey data work for DNK has required many more steps before its results could be translated into production garments and satisfied customers. First (and obviously), a body measurement is not a garment measurement. Skilful translation from measurements of bodies to *specifications* of dimensions of garments is necessary, particularly with respect to how garments 'grade up' from one size to the next. It is not

practical to calibrate every size option for every single garment for each season against the corpus of survey data. Accordingly, simplifying 'grade rules' need to be formulated for the 'grade ups' between garment sizes and the technologists at DNK have had to devise many new such rules. For example, while a size 12 standardly incremented to a size 14 by 5cm on the hips, a DNK grade may be in the order of 10cms. Devising these new rules is a considerable amount of work, still being reckoned with by the garment technologists in QC as we write.

In addition to the transfer of size survey results into product specifications for suppliers a series of *design recommendations* for buyers have been produced by QC. For example, one conclusion of the survey was that for certain groups of DNK's customers their hip measurements remained fairly constant but their waist measurements varied significantly. This result has led to an increase in the number of trousers being produced for these customers with elasticated waist bands. The company's replacement of outdated 1950s assumptions about population sizing with their own size data has also been reflected in the *models and mannequins* which are used to judge fit within the company. The company had two mannequins (or tailor's dummies) specially made for them on the basis of their new data and have recruited new people to work as in-house models. These people are employed to try on clothes for buyers or garment technologists and are now chosen for the correspondence of their body measurements with those of the statistical averages for the company's customers.

Getting Others to Get It Right: Design as Specification Work

While DNK's size survey has had profound implications for the way personnel within DNK organise their own work, because the company's products are not manufactured in-house, it was also vital to convince manufacturers about the new size specifications. The company is no longer ordering so many size 12s, so many size 14s and so on but is requiring that the garments be *their* size 12s, *their* size 14s. This required DNK's suppliers to change aspects of their manufacturing and to cut their patterns according to the new specifications and not their knowledge of the standard measurements established in the 1950s. These considerations suggest that fashion design in settings like the one we have studied needs to be understood in its inter-organizational context. Indeed, as one of the managers of QC told us when referring to his day-to-day involvement with the company's suppliers: "most of it is getting them to get it right".

Specification

Getting the others to get it right has become especially acute in the light of the size survey and DNK's consequent pursuit of ever better 'fit' for its customers. In particular, this has meant an increased importance on the ways in which the specifications of garments are formulated and presented to suppliers so that they can be acted upon. Let us look at the specification process in more depth. Typically a buyer will return from a shopping trip with a hundred or more garments and give them to a garment technologist together with instructions for the types of modifications that they want made to the design of each item. From this the technologist will produce a product specification for the modified garment design in the form of a size chart showing the ways key measurements grade up, together with a line-drawn 'sketch' of the garment. The sketch will be specifically drawn to highlight key features such as darts, embroidery locations and novel details which might otherwise be unexpected, together with identification of the points where the key measurements are to be taken.

The specification serves multiple organizational and inter-organizational purposes. First, a sufficiently skilled supplier should be able to, in principle, deliver items on the basis of the instructions it contains. Second, the specification defines what is 'right' for all practical purposes and against which samples will be judged once they are delivered. Third,

once a specification is agreed between DNK and a supplier, it may ultimately become a part of the legal agreement between those two parties. In short, the 'spec' is a translation from an existing garment and a set of handwritten comments into a series of requirements and sketches for a new design, which has a *productive*, *normative* and *legislative* standing.

Much, then, depends on the spec, and there are many hundreds to do for each catalogue in a very limited period of time. How do the technologists address these problems of 'time and scale' in their work? One of the important ways of attempting to save time which we have observed was for technologists to produce new specifications by modifying existing ones stored in the product specifications database. This *specification by modification* was extremely common and perhaps the most common way of producing a specification. Very few specifications were produced 'from scratch' but usually consisted of a measurement table and/or a diagram copied and modified from one or more existing ones in the light of the buyer's requirements. Indeed some technologists had specially produced a number of 'blanks' or template specifications which did not refer to specific products but which could be used as starting points for a whole class of garments such as summer dresses or sleeveless blouses and the like. The technologists use their skill to identify potentially similar specifications to act as a basis for modification into new designs.

Interestingly, the technologists often *collectively* identify specs, using *each other* as a resource for suggesting past examples which—in modified form—could meet current requirements. Indeed, this cooperative activity is a mundane feature of how 'specification by modification' is accomplished. Consider this excerpt from field notes:

```
Sarah is holding up a waistcoat and calls over to Jackie “Jackie have
you got any standard waistcoat measurements?” to which Jackie looks
over and replies “Done some - yer similar”. Sarah follows up with
"have you got any low returns on waistcoats?” and Jackie responds “no
they return high”. Sarah clarifies with “so you’ve nothing specific?”
but Jackie has been asked a question by another colleague. Sarah tries
again with “so I can take that?”. Jackie is still engaged in another
conversation and Sarah returns to the database. A few moments later
Sarah calls to Jackie again “Jackie I’ve got AA001, returning at 28
percent but 47 in [another company catalogue]”. Jackie looks up but
her telephone rings and she answers it. At which point two other tech-
nologists begin a conversation with Sarah as to how the same item can
have different returns rates in different catalogues. Sarah suggests
"basically the [catalogue] customer is not as fussy as the [other ca-
talogue] customer and not as happy about sending things back
(conversation continues)".
```

As well as indicating how the technologists have to cope with constant interruption, this excerpt shows how they use each other's knowledge of existing specs to help each other out. Furthermore, it indicates how design options are often considered in the light of returns data on similar garments and how these data are interpreted through collective judgement. In the current case, Sarah is able to make a reasoned judgement about whether to reuse Jackie's earlier specs informed by the returns data and Jackie's interpretation of them. On other occasions the technologists benefit from peripheral awareness of other conversations, searches or, even, the presence of a similar garment lying on the floor.

Jackie points to a top worn by someone who has just entered the room and says "that's CA444, fits nicely doesn’t it, I’m just doing a sub on that". Pat overhears this and asks “Jackie, so you’ve done a size spec for a top like that, do you know the number?” Jackie replies with “I think its on the front page” (of a catalogue)

It is important also to emphasise two other features of the technologists' work which help them collectively identify specs. First, it is very likely that—if one technologist is experiencing a problem or has a query to be answered—others have worked on the same range of items in the past, or are presently working on an essentially similar range, or are working on a range which has a few known differences. Second, the technologists share a room and this co-location affords the opportunity for peripheral monitoring of each other's work and ad hoc interaction occasioned by emerging problems. These features enable the technologists to sustain a working division of labour. While they each have their own specs to do, they can help each other out in suggesting possibilities and troubleshooting problems.

We have emphasised the multiple purposes of specifications and how specifications cross organizational boundaries. Garment specifications are *specifications for others*. As such they are designed with their recipients in mind. Ultimately, a spec will fall into the hands of a supplier who will act on its instructions. More specifically, it will be a pattern-cutter who will translate the size charts into the patterns used in production. Accordingly, one should not be surprised that the degree of detail and the exact form the specifications take often varies in relation to the specific supplier who will be its recipient or to the kind of garment being designed. We shall return to this point later where we shall discuss a garment that was believed to be over-specified.

Assessment

After buyers have found potential suppliers for garments, samples are sent to DNK for QC assessment. The decision which is made by garment technologists about a sample is recorded on a 'sealing report' which is sent to the supplier. In addition to passing or failing the sample, a section of the sealing report allows the technologist to write comments and recommendations. However, as well as this formal record of the assessment process and the exchange between suppliers and technologists of samples, technologists frequently liaise with suppliers about the details of samples using 'phone and fax. These calls often occur before the 'paperwork' is sent to suppliers and involve further explanation of what was considered wrong with a sample, what will be acceptable and in some cases suggested ways for manufacturing an acceptable garment.

Once a sample is agreed upon as being adequate to the spec, the design is said to be 'sealed' from further changes by either party—supplier or retailer—in that a supplier's production is supposed to correspond to the sample they have provided and the retailer has, on this basis, agreed the design and the make of the product and will not change their requirements. In these respects, sealed garments and sealing reports have a legal status which add to—and elucidate the interpretation of—the contract between DNK and the supplier itself.

There is a sense in which sealing also provides a 'closure' on the debates surrounding a garment's design and make. As well as specifying their design requirements the buyers and the technologists may have had to argue with suppliers for them too. Additionally within DNK the sealing of a design also seals agreement between the buyer and technologist concerned with the garment. Both have to sign sealing reports and this sometimes requires a reconciliation of different perspectives. Buyers may sometimes feel under pressure to get their selections and ideas translated into production, while technologists may be concerned that prematurely accepted samples may lead to high returns problems later on. This means that the adequacy of a sample or the accuracy of the supplier's interpretation of the spec are often a negotiated judgement produced by technologists and buyers working together.

That one should understand accuracy *in context* is also born out by the following example:

```
Jackie is still on the phone saying "... the only thing is the belt
loops - it is only a tiny weeny bit of fabric! - oh no wouldn't reject
it for that but if you could for the next delivery - it wasn't horren-
dous I could just see that with movement it was going to slip and
slide and the loops would help - ok then (transfers call) - yer the-
re's a call from Singapore for Jenny".
```

In this case a technologist is requesting that a supplier makes changes to the design of a garment—the addition of belt loops—but the supplier initially resists this suggestion on the grounds, we were told later, that the extra fabric would cost more. This is argued against by Jackie who points out that a belt loop is *"only a tiny weeny bit of fabric!"* The supplier is then interested to know whether QC would reject a garment if he did not add the loops and Jackie makes clear that—while it would not be adequate enough reason to fail a sample—she would still like to see the feature added in the next delivery. Here we see a technologist making a decision about what is worth rejecting supplies for and what is not. We can also appreciate the context within which such decisions take place: under time constraints, with the threat of holding up production (because "time is ticking on"), at the last minute when samples have been delivered late, in consideration of the many other product lines this supplier may make for DNK and so on. This context has consequences for the required degree of 'accuracy' in the correspondence of samples to their specifications. Accepting samples, as a step in obtaining production stocks, is a business decision not simply a technical one narrowly defined.

Investigation and Remedy

QC investigate the items within catalogues which are 'returning high'. The work involves both buyers and technologists meeting together and collaboratively trying to establish what is 'wrong' with a garment, why such problems arise, and how and if something can be done about it. Each garment under discussion will be viewed being worn by one or more of DNK's in-house models. The further analysis of high returning garments and resolving what should be done involves the juxtaposition of many different sources of information about the garment (e.g. its returns reasons, its sales 'performance') as well as considering its relations with similar designs in terns of 'parentage' (e.g. a different colour variant or variants which involve modification of the same 'ancestor') or history (e.g. another garment from a problematic supplier). The fidelity of production items to specification will also be tested by checking returned goods against the sealed sample to see if they are "measuring to chart". And—essentially—the model will be asked "how it feels".

At a high returns meeting one of us attended a coat with a returns rate of 39% was being extensively discussed. Most of the returns highlighted ill-fitting sleeves. This was an enigmatic—and for us revealing—case because the garment in production *did* accurately conform to specification. Julie, the garment technology manager in attendance, suggests that the spec includes "too may rules" through too many key measurement points being identified, and that they should consider giving the supplier "certain measurements and let them try and sort it out". Wendy, a garment technologist, observes that specification work is often done at a distance "sitting at a desk" and not in direct contact with the materials (the patterns and the fabric) which are used to actually produce the garment. As a result of over specification, the suppliers and particularly the pattern cutters "lose control over the silhouette" of garment features such as shoulder/sleeve joins just as they try to follow and reconcile a large number of key measurements. For a garment such as a coat and especially for a supplier who may be "trying too hard to please us", it can be "impossible"—according to Julie—to specify all details regarding "the insertion of the sleeves". Between them, Wendy, Julie, the buyer at the meeting, and a model wearing the troublesome coat, all contribute different observations on the problem and discuss a

number of possible remedies including: providing less specification (in this case fewer defined measurements but ones that "we do know"); providing specifications or guidelines of what they do *not* want ("saying this is what we are trying to avoid"); the direct modification of an actual physical garment (by "pinning it"); the use of graphical representations ("a page of shoulder shapes") which suggest critical lines or shapes rather than give actual measurements; and the use of photorealistic representations ("a picture").

This example and others from our field notes illustrate a number of points. *First,* the contribution of a number of different points of view (buyers, technologists, and—if necessary—their managers) are essential to the work done in investigating problems. Not only do the different attendees have different 'professional viewpoints', they also may have varying knowledge of the history of the specific garment under investigation. For example, one participant may know more about why the design and the supplier were selected and be able to bring information to the meeting beyond what is formally documented.

Second, a variety of remedies may be considered to address problematic cases in the light of the specific supplier involved, the garment in question, and the degree of return and demand associated with it. Indeed, difficulties with high return garments rarely lie in QC not understanding what changes need to be made. Very often it will be immediately obvious to 'the practised eye' what is wrong. Rather, what is required are means to get the supplier to implement those changes. For example, while it may be that, in general, size chart specifications together with the sketch they include serve as an adequate instruction for many garments and suppliers, if the case requires it, other means, other representations and—under some circumstances—even doing the pattern-cutting at DNK's own expense (or in-house or inviting the pattern cutter to a meeting at DNK to see the garments on mannequins or models) will be resorted to. This flexibility is essential to the smooth management of the supply chain.

Thirdly, the case of the ill-fitting coat suggests that it is quite possible for a garment to be *over-specified.* This is consistent with our remarks that accuracy has to be understood in context. In the work of QC, judgements of accuracy and degree of necessary detail are always made in the light of "getting others to get it right"—it is quite possible for excessive specification to slow the supply chain and jeopardise the prospects for timely delivery of correct(ed) garments or make it harder to manage.

A *final* observation to note about the fitting sessions when high return garments are being investigated is that the in-house models are full participants in them. Explicit requests for movements are very rare as instead models have the skill to know when and where to turn and what pose to adopt according to the on-going conversation to which they also contribute their opinions and knowledge. For example, if a discussion concerns bust darts the model may face front and square on to the other participants. Models do not seem to be discouraged from making contributions—after all they have the clothes on and there is no surrogate for their opinion as to whether it fits and how if feels. In addition they may well have tried on earlier samples, last year's lines, or other garments made by this manufacturer. They have an accumulated knowledge and skill which is 'embodied' in a quite literal sense and is invaluable to the investigation-work accomplished at returns meetings. Thus, while the correspondence of in-house models to tailors' dummies in terms of physical measurements is close, they are far from being dumb mannequins but are themselves skilled social actors in an occasion of cooperative work.

Garments and Representations

We remarked earlier that there are a number of different computer systems in use within QC each devoted to a specific purpose (product data, textile visualization, pattern production, returns analysis and so forth). Each of these technologies stores representations of (part of or an aspect of) the garment appropriate to the task at hand. If one is concerned

with the location of key measurements, these can be most readily depicted by means of a line diagram. If one is concerned with how different sizes grade up, then this can be most readily depicted in tabular form complemented by nested pattern blocks. If one is concerned to see how patterns can be cut to efficiently maximise the use of cloth, then a 2D projection of the pattern parts to show how they will lie on flattened cloth is called for. Much of the time it would be distracting to work with 3D visualizations of garments which combine all these details and give the whole garment a photorealistic rendering. Indeed, photorealism might 'hide' the idiosyncrasy of novel detail essential to the originality of the design and to communicating such to suppliers—matters mostly better conveyed by a sketch with the deliberate noting of such details and even their exaggeration. In fact, it is precisely details like seams and darts, which are so essential to both manufacture and the fit of the finished product, that often need to be most conspicuous in representations and least conspicuous in the garments themselves. That photorealistic representations might sometimes hide significant detail is only a paradox if one forgets what is obvious in the everyday work of garment technologists: representations are *recipient designed* and *purpose specific* for different aspects of design and troubleshooting supply chain problems (and accordingly have separate purpose specific technologies to help store and manipulate them).

This is not to say that there are never occasions when many aspects of a garment have to be seen together. Indeed, there are, and this can be especially vital in troubleshooting problems or rectifying faults. But there is a very effective way of seeing them all together: the garment itself! To paraphrase a well known slogan in current AI research: *a garment is its own best representation.* Production is practically *never* given the go ahead without a sample (this would only be considered if the exact same garment made by the same manufacturer was being reincluded in a later catalogue) and as soon as a sample arrives work can commence in the light of *that* to a much higher degree of confidence than if garment technologists were still manipulating representations no matter how veridical they were thought to be.

Furthermore, garments and the everyday objects associated with them (coathangers, clothes rails, pins) turn out to have important affordances for coordinating the work of QC. For example, the humble clothes rail is not merely a way of keeping clothes off the floor but also a way of *ordering* garments so that work complete (e.g. checked by a technologist on arrival from the supplier) can be separated from work pending at the other end of the rail. Work for the attention of one technologist is separated from that associated with another either by putting it on a different rail or by placing named cardboard separators between garments on a single rail. Related garments (e.g. from the same manufacturer) can have their coat hangers held together with a rubber band, thereby grouping them in an obvious physical way. Documentation and correspondence associated with a garment (e.g. specifications and sealing reports) can be placed in plastic wallets and pinned to the very garment they refer to. All of these features also make public and visible 'at a glance' the workload of the various technologists, its overall status, the progress made so far at this stage in the cycle of the seasons, and so forth. The garments themselves, then, afford cooperation and further aid in the technologists capitalising on their co-location in coordinating their work, helping each other out and maintaining a working division of labour.

Summary and Discussion

Let us summarise our findings so far. DNK Fashions has to coordinate the manufacturing of its garments by outsourced companies, which are distributed world-wide. This industrial and business context makes fashion design into an activity rather different from the more 'exotic' images of fashion design that one might intuitively have. Rather than design being conducted by isolated creative individuals who continually push the boun-

daries of taste and public acceptability (as an intuitive and somewhat 'romantic' notion of design might suggest), fashion design here has essential ties with its organisational and inter-organisational business setting.

Design is also tied to production and forever mindful of potential problems that an idea might encounter at the production phase. This makes a core component of fashion design 'specification work' and—like such work in other industrial contexts—specifications are rarely worked up from scratch. Fashion design is a matter of reuse just as much as many forms of software design. This means that design work is done as much with databases, records, sources of information and intelligence, as it is with artist's sketch pad or with CAD/CAM—indeed, more so. Furthermore, in the mass-market commercial context that DNK operates in, detailed knowledge of its customers and of the acceptability of current and past stock options (e.g. as revealed by returns analyses) become essential to reducing business risks and maximising the chances of satisfactory sales on items.

It is also vital to appreciate the socially distributed and variable nature of design activities. In fact, there is no single activity which organisation members refer to as 'design' and no single organisational locus for it. Design activity breaks down into 'buying trips', 'product specification', 'sample assessment', 'pattern production', 'textile design', 'sizing and grading', 'packaging and labelling' and so forth. All of these activities have their own distinct location within the organisation, their own specialised persons and supporting computer systems, and many are regarded as being 'technological' in nature. Furthermore, several of these activities can be contingently and variably located across organisational boundaries on a product-by-product basis and in response to emergent problems in supply chain management (e.g. pattern production is generally outsourced but the capability exists to perform it in-house if a pattern supplier has been unreliable and catalogue deadlines are approaching).

In contrast then to many intuitive notions of fashion design, in the setting we have studied, design manifests itself as a time-critical and seasonally variable activity, as a business concern with intimate ties with the whole manufacturing process, and as an activity involving the collection, manipulation and management of multiple forms of data about products, suppliers and customers.

Many key activities are essentially cooperative. While individuals may have their own distinct responsibilities, this comprises a working division of labour constituted by an awareness of each other's work and a readiness to help when required. Buyers and garment technologists cooperate on specifications, the sense given to 'accurately following them', what might be wrong with problematic garments or suppliers and what can be done to correct affairs. Technologists cooperate in assembling 'specs' and in retrieval of likely candidates for modification. This cooperation is essential to getting others to get it right.

The Uses of Ethnography

Over the two years we have had contact with DNK, we have been able to gain some insight into the nature of DNK as an organization, its relations with other organizations (in particular suppliers), the kinds of work done by buyers and technologists, quality control personnel and quality analysts, and been able to observe work as it happens. In the face of this detail and variation (from knowledge of interorganizational relations right through to cooperative work in real time), it would be trivialising to compress all that we have found out into one series of 'requirements' for cooperative technology. Besides we agree with [10] in arguing that ethnographic research should be an extra 'resource' for developers, a means of 'sensitising' developers to the nature of work in a setting relevant to technologies of interest (e.g. VR technologies within CSCW) and not something that can stand in lieu of traditional requirements documents and other resources routinely deployed in development. However, this does *not* mean that the only kind of intelligence that we

can offer is non-specific or vague. To serve as a resource is not to act vaguely or to be non-committal (contra [11] who construe an informing role as "reticence") . It is to appreciate that one's contribution is but one of many 'inputs' and it is to modestly recognise that one's knowledge of the contingencies of system development is most likely to be poorer than that of system developers themselves. The kind of study we have done—one which exposes us to a full range of organizational and interorganizational concerns—convinces us that there can be multiple 'uses of ethnography' in the development of CSCW systems and any attempt to demarcate between work studies on the basis of whether they can offer general or specific design guidelines will overly simplify the differences in kind that exist *within* the findings of most studies. By way of illustration of this argument, let us indicate some of these possible 'uses' and—in so doing—give a hint of the development projects we are assisting in.

(Re)shaping Agendas. Ethnographic work can be used as a resource for defining and redefining research agendas in CSCW. While many research problems emerge as significant within CSCW as a result of debates internal to CSCW itself or the disciplines which contribute to it, ethnography can help CSCW "learn from the field" [3]. Ethnographic work, by elucidating *members'* problems and concerns, can influence CSCW researchers in what we take as *our* problems and concerns. In the current case, this has led us to reappraise the possibilities for VR as a technology for CSCW in support of design.

We have emphasised that fashion design is no one single activity and bears scant resemblance to any abstract notion of what design might be all about. Accordingly, we see possibilities for collaborative VR systems in supporting cooperative information visualization and retrieval just as much as for systems which support creative individual 3D design. One of the endemic problems at DNK is coming to terms with the masses of information (on customers, returns, specifications etc.) required by and generated in the organization's work. Very often these sources of information are *juxtaposed* in problem solving and design work, and cooperatively so. As such, these activities fall within the remit of research on cooperative information visualization using VR technologies [e.g. 2]. And such applications of VR technology may turn out to be as relevant to design work as tools which support the manipulation of 3D splines or cloth animation. We are *not* saying that the advanced 3D design studio of our collaborators is misfounded, rather we are emphasising that there is a larger agenda for collaborative VR systems, for the activities which comprise fashion design work are many and varied.

Scoping Applicability. The results from an ethnographic study can often inform developers and implementors of the most appropriate settings for the technology they are developing or have developed. Such results may contradict intuitions about the usefulness of applications and an application may turn out to be useful in a setting not anticipated by its developers. This use of ethnography may be particularly prominent when development work is so far advanced or design decisions are so entrenched that it is impractical to change or reverse them. Consider, for example, the status of a 'virtual catwalk' in the light of our remarks above about the role that real models have in DNK's design work. We emphasised how DNK's models are valued not just for their shapes and sizes but also for their knowledge of the 'lines' which the catalogues have contained over the years, their abilities to orient to technologists' and buyers' interests in garments without them having to be explicitly told to adopt a certain posture, and their readiness to inform on 'how it feels'. This is a long way away from the most read-to-hand connotations of the 'virtual catwalk' (e.g. where virtual models parade garments as part of a fashion *show*) and what is most readily developed as such a technology. Accordingly, we should not think of virtual mannequins as substitutes for real models or of the virtual catwalk as offering support for the same activities as real models wearing real clothes do (e.g. investigation-work).

Similarly, the concept of an advanced 3D design studio for fashion may seem to be 'squeezed' in its applicability by our remarks about the standing of representations in design work. Separate, purpose specific and typically 2D, textual or tabular representations seem to be most appropriate at the early stages of design and shortly thereafter a sample of the garment itself will appear. Where—between initial conception (which typically is itself re-use) and a garment (its own best representation remember)—is the place for advanced 3D and photorealistic visualization? Of course, there are several places but they too are associated with specific problems. For example, to accurately specify some garments, some key measurements need to be taken from the back to the front, so the two 2D line drawings (front and back) which depict their location on specification charts do not convey the locations especially well. It is known that suppliers often find this confusing. Furthermore, cloth with large, bold textile patterns (e.g. a large check) often needs to be cut carefully so that the pattern coordinates with the seams, darts and folds of a garment. Again, such a requirement underwrites the usefulness of 3D representation of a more photorealistic sort than can be typically achieved with basic line drawings. Whatever the purposes advanced 3D garment visualization and animation need meet, they are *just as specific* as the purposes that the more commonly used 2D line drawings and tables fulfil. One should not have the notion nor develop design support technologies with the view that line drawings and the rest are somehow inferior because they are less veridical or are 'dimensionally challenged'. In a context of design-for-manufacture, they may often be just the job much as 3D visualizations and animations shown on a high power lap-top to a supplier during a buying trip may be appropriate other times.

Finally with respect to notions of an advanced 3D design studio, it is important to emphasise that design work at DNK is done with many forms of organizationally relevant information at hand. Designing and visualising a garment is never a matter of 'purely' appreciating its aesthetics. The selection of specifications is done in the light of sales figures, returns analyses, the performance of suppliers and so forth. Whatever else a virtual design studio might be, it should provide access to other forms of organizationally relevant information as required. Accordingly, we are helping explore not merely the development of cooperative information visualization applications but also their integration with design tools.

Beyond Interface Metaphors: The Virtual Objects of Work. Through a detailed study of the workplace, ethnographic research can uncover many of the means by which cooperative work is coordinated by members. This can often be suggestive of overall metaphors which can guide application development and present intelligible interface objects. However, much of the excitement surrounding VR as a technology for CSCW lies not just in its ability to provide 3D interfaces so much as in the intention to render *spaces* which could serve as *environments for cooperative work* involving the manipulation of *virtual objects*. Of the artefacts to be found in the real-world spaces at DNK, it is perhaps those which afford the coordination of work which are of most interest to the developers of collaborative VR systems. Accordingly, we are helping explore the utility of such mundane objects as rubber bands, coathangers and clothes rails for translation into virtual counterparts which might aid the coordination of work in virtual spaces. In addition to objects of work of these sorts, we are exploring the possibility of a 'Virtual Trade Show' as a space to present information about suppliers in the fashion industry.

Informing Development Choices. When ethnographic work is being conducted concurrently with development work, a detailed knowledge of a relevant field site can influence specific design choices within a given application. For example, consider the product database in use at DNK. This application does support 'specification by modification' but only to a limited extent. It supports the copying and pasting of data from one entry to another and the use of clip art. However, this is a long way from designing a system from the ground up that is intended to support 'specification by modification' as it is done at DNK. Enthusiasts of object-oriented database technology will obviously see possibili-

ties here but also some more detailed features of the specification process could be reflected in a revised database. For example, as we have discussed, there is an important distinction between specifications that are 'sealed' upon an acceptable sample being received and those which are not. However, it is *not* possible after sealing a sample to prevent changes in the electronic version of the specification and it is a paper print out which accompanies the sealed sample to the warehouse and is retained as a definitive copy by QC. This means that the database and the binding specifications can sometimes (inadvertently) get out of sync and that there is no directly supported way to do searches on sealed versus non-sealed specifications. Furthermore, the database is essentially an extension of a single user application. As a result it gives little support for the mutual awareness users might require of each others' activities. Users—even those concurrently accessing the same entry—are 'insulated' from each other. Clearly, from our account of cooperative work at DNK, it would be justified to pursue shared database technology as of relevance to fashion specification work. Thus, within a remit to develop database applications, we feel that our field work can be a source of quite specific intelligence.

Contextualising Evaluation. Furthermore (and finally on *this* list of 'uses of ethnography'), we feel our field study can inform the evaluation of emerging CSCW systems—be they virtual catwalks or shared product-databases. Our account of fashion design work at DNK highlights issues of time and scale, and organizational and interorganizational coordination. This provides a context for the assessment of any system intended for supporting work in the fashion industry. How does the system promote the speedy production of design alternatives? How does the system save time in the cycle from initial specification to starting production? How does the system enable the many hundred/thousands/millions of representations to be made organizationally relevant? Are the ways in which it combines data appropriate? Does it manage the data and keep it ready-to-hand in just the right way for deployment in cooperative work settings? How does the system facilitate the management of the supply chain? How does the system help getting others to get it right?

Ethnography in System Development Revisited. Let us close by returning to the issue of the relation between ethnographic research and system development in CSCW and first remind the reader of the context of our research being conducted as it is within a joint industrial/academic project involving multiple partners. Some of the other partners we had worked with before, some not. Some had ongoing research programmes which the project itself just formed a part of (and which—therefore—we would scarcely be entitled to overturn even if we wanted to), others were interested in finding out things afresh. Some partners provided field sites for ethnographic work in the project but other fieldwork (such as ours) was located in a company who participated purely through goodwill. This variability within the project and the many different kinds of thing we have found out about fashion design work convince us that it is wrong to excessively schematise the ethno-techno link.

We have already indicated this by noting the many different uses we think different parts of our work can be put to. Let us continue by revisiting Hughes and his colleagues [9] and their four different 'roles' for ethnography in system design: 'concurrent', 'quick and dirty' and 'evaluative' ethnographies, and the 'reassessment of previous studies'. Principally, the four ethnographies are distinguished from each other in terms of where they fit into the 'design cycle' (i.e. concurrently and along the same time course as development work, or initially and quickly to get things started, or once prototypes have been developed or systems released, and so forth). While elements of Hughes et al's different 'roles' are true to our experience, we suggest that the degree to which development agendas can be changed or are resistant (and at what level of detail) are further factors of note. Indeed, it is these considerations which give the time course of development its pertinence as a constraint (or otherwise) on the impacts and form of field research.

In our case, we would not be able (even if we had wanted to—and we don't) to halt our collaborators' work on the virtual catwalk and the 3D design studio. They were committed to work on these technologies from day one. What we can do is scope the applicability of them, suggest virtual objects to include in them to support the coordination of work, and indicate the importance of integrating them with sources of organizationally relevant information. That is, although we could be regarded as practising concurrent ethnography, not all features of these applications could be regarded as co-developed. In contrast, large scale information management has been a persistent theme of our work and indicated—quite early on—to some of our other collaborators (and us) possible uses in fashion design of information visualization techniques with collaborative VR technology. However, neither they nor us could spare the resources to further develop these ideas until fairly recently. Thus, what could have been concurrent ethno-techno development might—unless we are able to resource further fieldwork—turn into development done in the light of a reassessment of a previously completed study.

While Hughes and colleagues have done a valuable service in emphasising that there are not one but several different ways in which ethnographic research can contribute to development, we now feel there is a danger that their distinctions could be reified in the hands of others into *the* four ethnographic methods for CSCW ([9] though is quite clear in warning against this). Similarly, there is a danger of reading Plowman and her colleagues [11, especially p.319-320] as classifying different studies as if each had only one type of outcome, relevant to just one phase of system design. However, the contingencies of system design—much like garment design—exceed simple depictions. If 'getting others to get it right' is ad hoc-ed, negotiated and variably accomplished when tables of numbers are exchanged, why should it be otherwise when ethnographic reports are being passed around and who's 'supplying' whom is not always settled?

Acknowledgements

We thank everyone at DNK Fashions for their help and tolerance, and Steve Gray and Tony Rosella for their insights into the fashion industry.

References

1. Benford, S. D., Bowers, J., Gray, S., Leevers, D., Rodden, T., Rogers, A. & Stanger, V. 1994. The VIRTUOSI project. *Proc. VR'94*.

2. Benford, S., Snowdon, D., Greenhalgh, C., Ingram, R., Knox, I. & Brown, C. VR-VIBE: A Virtual Environment for Co-operative Information Retrieval, *Proc. Eurographics'95,* 1995.

3. Bowers, J. (1995). Making it work. *Information Society*, 11, 189-207.

4. Bowers, J., Button, G. & Sharrock, W. Workflow from Within and Without. *Proc. ECSCW95,* 1995.

5. Bowers, J., O'Brien, J. & Pycock, J. Practically accomplishing immersion. In these proceedings, 1996.

6. Bowers, J., Pycock, J. & O'Brien, J. Talk and Embodiment in Collaborative Virtual Environments, *Proc. CHI'96,* 1996.

7. Greenhalgh, C. and Benford, S. MASSIVE: A Virtual Reality System for Tele-conferencing, to appear in *ACM: TOCHI.*

8. Heath, C. & Luff, P. Collaboration and control. *CSCW*, 1 (1992), 69-94.

9. Hughes, J., King, V., Rodden, T. & Andersen, H. 1994. Moving out of the control room. *Proc. CSCW'94*.

10. Hughes, J., Randall, D. & Shapiro, D. 1992. Faltering from ethnography to design. *Proc. CSCW'92*.

11. Plowman, L., Rogers, Y. & Ramage, M. 1995. What are workplace studies for? In *Proc. ECSCW95*.

12. Randall, D., Rouncefield, M. & Hughes, J. 1995. Chalk and cheese. *Proc. ECSCW95*.

13. Rouncefield, M., Hughes, J., Rodden, T. & Viller, S. 1994. Working with "constant interruption". *Proc. CSCW'94*.

14. Shapiro, D. 1994. The limits of ethnography. *Proc. CSCW'94*.