

CSCW AND UBICOMP INTRO

Stuart Reeves – stuartr@dcs.gla.ac.uk

<http://www.dcs.gla.ac.uk/~stuartr>

1 Computer supported cooperative work

Computer supported cooperative work (CSCW) is generally concerned with two things:

- Studying groups of users (perhaps using technology, or perhaps not) and analysing and evaluating their 'work' practices
- Designing technology - using the results of evaluation to inform good design

(Although we use the term 'work' for CSCW, it is important to note that use of this has typically been an analytic tool rather than meaning 'workplace work'. So, 'work' here is a generic conception of the word, e.g., what is the 'work' of organising a leisure activity with friends, what is the 'work' of playing a video game.)

CSCW has an important relationship to the larger field of Human Computer Interaction (HCI). Within HCI it is fair to say that historically human-machine interaction has often been conceived of as a *single user in dialogue with a machine*.

In the 80s researchers realised that they were neglecting the reality of these situations. Instead they decided that interaction may still be considered as a dialogue, however the interface is bigger than we imagine. Researchers were also beginning to examine settings outside the laboratory, such as workplaces.

CSCW became concerned with this greater perspective on the interface's scope, and also began to examine real interaction with machines found in the world, rather than the idealised interaction often found within laboratory experiments.

Bentley et al. in their paper on air traffic control summarise the conceptual shift that was taking place at the time:

"The notion of the user as a stable generic entity with fixed cognitive capacities is a chimera, and a model which cannot survive the move toward larger scale, distributed systems."

2 Problems with HCI

Here we can start to explore some of the problems with HCI. As mentioned above, one major problem was the strong focus on the individual. Also:

- Historically HCI has not dealt with multi-user situations, e.g., networked environments such as virtual worlds.
- There has been a reliance on experimentation in controlled environment versus testing things 'in the wild' or in situ.
- HCI hasn't usually taken into account what people actually do as part of their work—it has ignored the contextually relevant aspects of use of an interface in a particular location.
- HCI typically looks at performative measures e.g., how fast people type text messages, but doesn't say what are text messages are actually used for?

- Evaluations are typically too short—they focus on the first 5 minutes of interaction, not the next 10 years (to quote Buxton [REF]).
- Ignored more exotic applications, e.g., entertainment, performance and games

CSCW came and impacted this picture, changing and influencing HCI.

At the time this was set against a backdrop of computers increasingly being used by people in groups—the individual was no longer the user; computers increasingly networked via LANs; workplaces were becoming full of computers and networks of computers; even non-networked computers used in large groups; the rise of applications such as email, file sharing, message forums in late 1970s; the rise of the internet.

This perspective in HCI led indirectly to various systems failures. These were not because of the interface—HCI hadn't failed in one sense. The systems failures were to do with social and organisational issues that weren't taken into account—systems frequently didn't fit with how things were really done and so were bound to fail at some point. The key question became, how does computing fit into social and organisational work practices? This leads to CSCW.

This problem was important and still is important.

In 1992 the London ambulance service had a new computer aided dispatch system to assist the scheduling of ambulances and so on. It was thought that digitising the system would increase efficiency and so on. But...

- When rolled out the system's use lead to huge delays
- Around 30 people may have died as a result of these

The report said:

"Management were misguided or naive in believing that computer systems in themselves could bring about changes in human practice. [...] They only assist in the process and any attempt to force change through the introduction of a system with the characteristics of an operational 'straightjacket' would be potentially doomed to failure"

(London Ambulance Service Inquiry Report, Page et al. 1993:40)

That was 1992—but there are still problems. For example, the Sainsbury's Accenture flop:

http://www.theregister.co.uk/2004/10/19/sainsburys_v_accenture/

As a way of sensitising ourselves to how the HCI problems might occur let's consider the how we might digitise medical records. On the face of it we should just be able to recreate case records using as a series of fields in a database, e.g., name, date of admission, etc. But the use of the records in practice reveals that this is too simplistic.

We might find from a study:


- Doctors can assess at-a-glance the age and frequency of use from the visible wear and tear on the record.
- Handwritten notes give clues as to the number of different individuals updating the record.
- The way the form is arranged spatially informs how it is to be read—something potentially lost in arbitrary data storage systems.
- There is only one physical record potentially—its presence or absence from a location means a particular thing—something we lose with digital records.

So the record is generally very rich and the design of our system to replace medical records should take into account these things—the practical, mundane use of such objects in situ. Straight digitisation misses the often subtle ways things such as records are used.

3 Changing the perspective

One useful way to view this changing perspective is Jonathan Grudin's 1990 paper on the interface. The diagram has a user, but their interface consists not just of the interface software and I/O devices but also system administrators, documentation, colleagues and so on.

Interaction is more like being part of constellations rather than a dyad.



This isn't the full story though: it is not just workplaces that are relevant, computation is also spreading into public settings like museums and galleries, theatres, clubs perhaps as part of performance art, mobile games, music performance, or some educational purpose.


If we begin to engage with this wider understanding of the interface and its spread, then some of the following becomes necessary:

- We must start to study interaction in its original contexts so as to retain access to these environmental and infrastructure aspects of the 'interface' that we might otherwise lose in the move to the lab.
- Conceptually we must shift our perspective to locate frequently examined parts of these constellations (i.e., user and machine) and analyse them in the context of these constellations.

- We also must begin to consider other parts of these constellations that have often been overlooked (e.g., the role of colleagues in the production of an individual's work practice at the computer).

4 CSCW systems

For a more broad overview, this diagram (from wikipedia!) gives us some rough idea of what kinds of software CSCW is interested in studying and designing.



This is an early CSCW framework which derived from studies of 'groupware'. Groupware applications were a major CSCW technology that was developed and built upon in order to explore CSCW concepts and demonstrate collaborative work issues and solutions. Often these are basic collaborative tasks that most of us should be familiar with nowadays, e.g., messenger style applications, email, single shared displays, etc. CSCW research pioneered many of these systems.

The framework provides one way of thinking about the different kinds of applications we have:

- *Synchronous activities*, technologies that help people work at the same time (usually in different places), and are 'real-time', e.g., virtual environments, video conferencing
- *Asynchronous activities*, technologies that help people work at different times (usually in different places), and are 'non-real-time', have an acceptable latency etc. e.g., email, discussion forums, databases
- *Distributed activities*, in different locations, e.g., virtual environments
- *Co-located activities*, taking places in the same or similar physical space, e.g., working together around a computer screen

5 Core CSCW topics

Some of the core concepts that are of relevance to CSCW work—they are all interlinked as the studies will show when we look at them, but there are a few we can tease out (note that it's important not to see these as discrete—many of the studies presented in this document share these concepts).

5.1 Collaboration with, around, and through artefacts

The first topic is the mechanisms of collaboration which involves collaboration with, around and through artefacts. Lots of CSCW studies have looked at how people collaborate with, around and through artefacts, both technological artefacts and non-technological ones.

Paper is a particularly good example. We can think of the vast number of ways paper can be used to organise everyday life. Post-it notes, bus timetables, books, forms, instruction manuals, etc.

Collaboration often takes place through artefacts as well as through talk, e.g., the medical records example outlined previously above. It can also take less tangible forms: dictionaries, plans, rulebooks, categorisation schemes, classifications, web forms, fare rules, regulations, procedures, traditions, instructions, agreements, tacit acknowledgements, physical arrangements of buildings, furniture and other objects. People collaborate around objects like these every day.

In our work this increasingly involves collaboration and coordination around technological artefacts. In order to design technology sensitively we have to understand a number of things.

We can start by asking the following questions:

- Firstly, how do people coordinate their activities around non-computational artefacts? What are the competencies they bring to bear upon that coordination?
- Secondly, how do their coordinative activities change when we introduce technology into the situation?


5.1.1 Air Traffic Control study

The first study to look at is one involving air traffic control from 1995. Harper and Hughes studied how air traffic controllers manage air space using 'flight strips'. As Dourish in his book "Where the action is" describes:

"the topic for this investigation was not the rules and procedures of air traffic control as they might be laid down in an organizational manual. Instead, it looked at the actual practice of air traffic control as it occurs and unfolds, moment by moment, day to day, and as it is executed and experienced by the air traffic controllers themselves."

In the study there were two controllers, two assistants and one sector chief (collectively called a 'suite'). The suite is generally responsible for one sector of airspace. They use radar and comms equipment to manage that airspace.

A flight strip represents a plane in the air through a particular sector, and shows information like call sign of the aircraft, speed, altitude, direction, etc. (see below).



These strips are put together on a board (see below).



The strips are the focus of the activities of the controllers and how they manage the airspace even though everything is available digitally.

Controllers manage the airspace by physically manipulating and managing the strips. Strips can be annotated, folded, moved up and down racks. The physical configuration—i.e., where and how the strips are organised spatially—indicates things like the status of the flights, whether someone wants attention drawn to them by pulling them out of alignment (see figure again). Physically passing the strip on indicates the plane moving from one space (and one controller) to another.

Strips also represent the activity of the controller managing the space, and are readily available at-a-glance for other members of the suite team. They help the team distribute their activities and coordinate with one another. New controllers arriving for the next shift can spectate on the existing ones' manipulations of the strips and get a sense of the current state of the airspace they are about to manage.

The study revealed details of how air traffic controllers did their job—how the computer system mostly augmented an already existing paper-based activity. So we have to think carefully before we attempt to digitise systems like this and introduce technological versions of existing activities. For instance, a new digital system for managing airspaces might not support many of the vital features that the study uncovers. For example, a digital display can't necessarily be 'folded' (unless we design in some digital way to do so), and it can't necessarily have arbitrary annotations on it (again, unless we realise it is important to devise a method to incorporate them).

Thus, we can use observations like these to drive better design.

5.2 Flow of work

This second topic is about commonsense notions of rules and procedures in a flow of work compared with the work that people actually, really engage in to achieve something. What we might call 'business as usual' accounts gloss the range of


activities that are engaged in which are pertinent to design (as we saw with the medical records example). In some sense you can't trust what people say they do at work—what they really do is different from these glossed accounts they give.

As Schmidt and Bannon state of this topic:

“According to the traditional ‘bureaucratic’ conception of organizational work, people perform a number of tasks according to a set of well-specified ‘procedures’ that have been developed by management as efficient and effective means to certain ends. The traditional formal organization chart is presumed to show the actual lines of authority and the correct’ pattern of information flow and communication. However, the conception has been proved highly idealized and grossly inadequate for analyzing and modelling the articulation of real world cooperative work arrangements.”

5.2.1 Reported flow of work

The diagram here is from Paul Dourish's book “Where the action is” and depicts a classical HCI perspective on how this flow of work is achieved.



In this model the plan divides up the task into achievable goals, sub goals and operations in order to reach those sub goals and eventually goals. There is an ‘execution cycle’ engaged in by users who perform operations, compare the results of those operations via monitoring to the plan as it exists and its various sub goals

Donald Norman pointed out the ‘gulf of execution’ present in this view. In the interactional dialogue, a user and designer’s “conceptual models” of the functionality the interface provides access to are often found conflicting. This results in “gulfs” of execution and evaluation where the user struggles to successfully predict the outcome of an action.

“Plans and Situated Actions” is all about the problems of human-machine communication and this planning perspective. It was published in 1987 and is by Lucy Suchman, an anthropologist. She studied how work actually took place with photocopiers as her subject. She used ethnographic techniques to achieve this.

Suchman discovered that this idea of ‘work flow’, i.e., a deconstructed series of plans in order to achieve some goal was something of a myth. Instead, people rationalise, post-hoc, about how their work happens in order to make it easier to understand and more acceptably accountable to others. Very little work is as routinised as plans and goals would presume—most work involves sidelines, exceptions, and so often that ‘exceptions are the rule’.

The book examined how interaction requires mutual intelligibility or shared understanding: key ethnomethodological topics (which we’ll discuss later).

The book was set against the cognitive science, task modelling and AI of the time. This perspective modelled the planning / goals view of human activities and

implemented this in machines (as we saw above), but ignored aspects of how plans / goals are used in the real world and how mutual intelligibility is established and maintained.

The book also pushed against the view that intelligence is symbol manipulation, that the mind is a computer, that mental states are simple logical statements, that activity can be modelled as a set of simple rules or speech acts.

5.2.2 Plans and situated actions

In order to see the actual flow of work at photocopiers Suchman divided her data (video) up into four strips of action (see figure): what was not available to the machine on the users' side; what was available to the machine on the users' side; what was available to the users on the machine's side; and the design rationale (i.e., the intent) behind the machine's displays.

THE USERS		THE MACHINE	
Not available to the machine	Available to the machine	Available to the user	Design rationale
B: (Hands on latch) A: "And lift up to the left." (Looks to B, who struggles with the latch) "Lift up and to the left." B: (Still struggling) A: Okay:: B: Pull, and lift up to the left. (Looks at picture) Oh, the whole thing A: [Yea. B: Lift up and to the left.			
RAISES DOCUMENT HANDLER		DISPLAY 3	Instructions for placing document and closing document cover.
A: "Place your original face down (Passes journal to B) on the glass, centered over the registration guide." (Looks to machine) Got that? (pause) Want me to read it again? B: Um:: I'm just trying to figure out what a registration guide is, but I guess that's this, um:			

Simply doing this presentation of the data reveals the two threads—the human side and the machine side. It reveals that there is the problem of the system's opacity and autonomy. Thus, you can't see inside the machine, and the internal mechanism is complex and to some degree independent of you.

Also, the photocopier design had various in-built assumptions: the system is designed to follow a script; the script prescribes and controls actions that the user may engage in.

With a script-like perspective on interaction, the design is goal-directed, planned and proceduralised—i.e., following a reported flow of work instead of an actual one. In this design all possible goals, sub goals and operations are defined by the system script. For each step in the interaction, the system checks the adequacy of each stage of the procedure and acts accordingly. So, in a script a basic procedure of interaction might go something like this:

- System presents instruction
- User reads instruction, interprets references does the action
- Design assumes that this action means that the user has understood the instruction
- System presents next instruction
- And so on...

But this system is inadequate. We get breakdowns of interaction like 'false alarms' where the user misinterprets something presented to them by the machine. In this situation the machine incorrectly deduces that something is wrong, then stops or does something internally (i.e., 'thrashes'). This is a symptom of system's opacity, i.e., its lack of good feedback and not presenting its model of activity to the user.

The system is also unable to successfully relate the relevance of past failed actions to the current action. This is illustrated by the 'garden path' problem. Here a user is uncertain of what the procedure is, their action is wrong but it fits with current requirements of the system to elicit the next response. The user(s) carry on and make sense of subsequent actions in the local context and interpret it to fit as best they can with what just happened—documentary method of interpretation (see Garfinkel later). System continues on too, but the original error's effect appears later... however the origin or history of breakdown is no longer easy to relate to the error the user sees.

The essential problem then is a mismatch between the users' actual activities—i.e., including all that articulation work—and the system's model of users' activities (which is incorrect).

This is not to say that plans, goals, sub-goals and so on are nonsense. They have just been misinterpreted. The key point then is that the plan is a resource to draw on, not a script to follow. Suchman gives the example of the kayaker in order to highlight this:

- They have a rough idea of where they wish to go downstream however they cannot plan explicitly.
- They have to react in situ to the moment-by-moment changes in character of the water.
- They improvise in order to be more flexible than their plan alone.

5.3 Awareness

Social interaction with and around technology often involves many levels of awareness around the system, which are not necessarily accessible to it. As we saw with the Grudin diagram, the interface is not just the keyboard of the machine sitting in front of you. So awareness as a design issue is bigger than just input and output.

This implies a better model of awareness that takes this into account, one which appreciates that awareness of one form or another is central for many day-to-day activities.

Consider how awareness of one another is understood in a non-technical world: talk, observation of bodily conduct, and general seeing. Then consider how your awareness of one another is hugely modified when we introduce technology: the internet, mobile phone networks, etc.

One psychological way we can think about what awareness is follows:

- We can consider direct sensory inputs: light, sound, etc. which can then be broken down into 'objective' phenomena e.g. wavelengths.
- We can then think about the construction of a symbol to classify input state, e.g., if some light is in a particular range of wavelengths then it 'means' red.
- Once we have done this we can fit that symbol into goal-directed algorithm or plan, e.g., if I approach a red light then I stop my car.

But this doesn't fit with observations since real-life classification is far more dynamic and contextual. Further, this kind of plan-based action is too simple a model of most of everyday life, remember what was discussed with regard to Suchman's study.

Awareness is also social; if we model awareness based purely in psychological awareness we'll miss that out. Instead we can try to understand how awareness is mutually established, in a negotiated way between people.

An important line of CSCW work has been examining workplaces such as control rooms, e.g., ATC stuff we saw previously. We can also see how awareness and coordination has been addressed in these settings. So, with the ATC study, where we looked at the objects that are used for collaboration, here we will look at how awareness, e.g., peripheral awareness, and corresponding mutual coordination through continual monitoring (i.e., watching others) comes into play.

5.3.1 London Underground study

This section looks at Luff and Heath's study of London underground control rooms. In this study, the Line Controller coordinates the day to day running of the railway. The Divisional Information Assistant (DIA—i.e., announcer) is also present, whose responsibilities include providing information to passengers through a PA system and communicating with station managers (see figure).



Mutual awareness of what was going on among the staff in the control room was key in this study. When there was a problem on one of the lines, local staff overheard, and started action that took into account that problem. Awareness supported the staff working together. We'll look at each of the topics raised by the paper in turn.

Surreptitious monitoring and interrelating tasks

In this situation the DIA makes an announcement to passengers based on overhearing the controller's conversation with the driver and assessing its implications for the expectations and experience of travellers using the service. The DIA transforms the controllers request into a relevant announcement by determining who the decision will affect and its consequences.

Rendering activities visible

Here the controller talks through the timetable out loud thus rendering what is a 'private' activity publicly visible. This avoids establishing explicit mutual engagement with colleagues which would undermine the ongoing accomplishment of the task in question. But permits a way for the DIA to use the talk out loud as a resource.

Overseeing the local environment of events and activities

The flow of information and the responsibility to oversee the environment for the other is not simply one-way. The Controller assumes responsibility for keeping his colleague informed of events which may otherwise pass unnoticed. The DIA will monitor the operation of the service and draw his colleagues attention to any events or problems which may have been missed.

5.4 Appropriation and adaptation of technology

This is about users developing competence over time. It has particular relevance for what are known as ubicomp systems—i.e., ubiquitous computing systems.

Following lectures will present more information on ubicomp systems, however for our purposes it covers all sorts of mobile devices people use.

These necessarily involve adaptation and appropriation because they are used over long periods of time and are typically design explicitly to be woven into the fabric of everyday life. (Although the same could be partially said about desktop computer systems—they are appropriated and adapted to but this was not the explicit design intention.)


5.4.1 Can You See Me Now? study

A mobile game called 'Can You See Me Now?' (CYSMN) demonstrates aspects of this appropriation and adaptation. It also demonstrates a moving beyond understanding interaction in terms of the first 5 seconds (i.e., because it looks at long term development of skillsets).

CYSMN is a simple game of catch played out across real and virtual environments. Professional performers are runners on the streets of a city, equipped with a GPS and a PDA. On their PDA they can see the positions of players. Players (members of the public) play online on the web, moving their virtual character on a map. Runners' positions are also displayed on a map. When runners get within a few metres of the virtual, online player, the player is caught. The game is played over the course of many days with hundreds of online players

Of interest is the way GPS is used to mediate location between the real and the virtual. This means understanding how runners adapted to and appropriated the vagaries of GPS technology. In order to assess this an ethnographic study shadowing runners' activities was done.

Runners developed an understanding of GPS shadows—i.e., places in which GPS satellites had problems covering. The diagram shows GPS shadows in Glasgow, at 17:00 and 17:05. Grey shadows show areas where GPS coverage is likely to be poor.



The runners then exploited these shadows within the game, in order to surprise players. They developing an understanding of good times of day to go certain places.

Essentially, the runners started to develop a stock of knowledge (expertise) for interaction mediated by GPS, i.e., GPS is the medium through which runners and online players interact. As part of this the team met every day to discuss tactics and strategies. These tactics and strategies emerged over time. Through the meetings the runners had adapted to GPS technology and appropriated it as part of their performance work.

6 Evaluation techniques

The previous sections looked at a number of studies. But how do we go about finding these kinds of things out in the first place? What techniques can we use

which help us take into account the complex and rich settings which technology is designed for and deployed in?

Many of the studies that were presented earlier on were ethnographic in character. The subtleties of their insights were developed from careful ethnography. These were in depth investigations where the researchers went into the workplace to examine what people actually did, for example, what people actually do rather than what managers said they did.

Ethnography was a radically different approach to studying technology use when people first started using it in CSCW. It has been massively influential in CSCW; it helped get around some of the considerable problems facing HCI when addressing multi-party, collaborative settings. Ethnography has also become mainstream—big companies like Intel and IBM have a crew of anthropologists who use ethnography in order to help design their systems better. Ethnography is non-experimental, and it's done outside of the laboratory. The settings are not controlled, they are explored as-is, in situ. This means ethnography involves studying systems used in their 'natural' habitat, rather like how a biologist might go out into the environment and study animal behaviour.

Ethnography is the process of trying to understand the social world through the eyes of its participants. Malinowski said it was "to grasp the native's point of view, his relation to life, to realize his vision of his world." With ethnography we have an advantage though, since humans are able to make sense of one another's activities with quite a lot of certainty thanks to being human, so in that sense it is not like studying animal behaviour.

Ethnography's roots are in anthropology. Historically in anthropology, researchers who wished to study and understand particular tribes at first relied on books, documents, and artefacts from that particular culture, and did their anthropology using this. Later anthropologists pushed this further and often lived with that tribe for a prolonged period of time. This technique is ethnographic. Malinowski was one of the first anthropologists who used ethnography. He helped introduce an important technique in ethnography called participant observation which involves observing and potentially taking part in the group under study.

There are varying levels of participation that is possible; obviously for things like ATC, it was hard for ethnographers to take part without spending years training, so they just observed rather than participated. For other studies, on the other hand, ethnographers may start taking part in the activity which is of interest, in order to help them see how it is done. In this way the researcher may become a participant in the group, not an external observer looking in from the outside. It varies between studies depending upon practicalities, like time, resources, etc.

As part of their ethnography, the ethnographer will also take copious notes on their observations and participation, and they will collect other artefacts that are relevant, e.g., copies of documents, etc. This might also involve photographs, audio recordings, video recordings, etc. The ethnographer might use techniques like shadowing a person for a day, following an object through some process (e.g., some object people are working on such as something being assembled), observing a particular setting, e.g., a reception desk, or observing a particular activity or process from start to finish, e.g., going out on call with ambulances.

Of course, ethnographers also do interviews too; these can be more or less useful documents for different kinds of ethnography. Typically interviews are quite open ended and free of much structure. Further to this, different kinds of ethnography might place more or less emphasis on things like video data, interview data and so on.

Regardless of the precise level of participation or observation, the ethnographer begins to organise what they have observed and participated in into a rich

description, perhaps organising things into key themes. The result is an ethnographic report which documents the 'life world' of the setting under study. Its key themes will hopefully bring out the important practices which members of the setting engage in—i.e., the practices which make that particular setting just the way it is

7 Ethnographic forms

There are various underlying theoretic principles ethnographers orient to when constructing their reports; they may have a variety of motivations for producing their reports the way they do. Different theoretic and / or philosophical principles essentially provide some techniques for ethnographers to structure their findings: deciding what is important and what is not; what is the focus of their report; how to interpret what they have observed; often their whole approach to the topic.

Having philosophically different groundings informs what an ethnographer thinks they are able to make sense of in their observations. We'll cover some of them briefly here and go on to discuss in greater detail one of the most influential in CSCW (and HCI more broadly): ethnomethodologically-informed ethnography. (Note that it is worth reading the Sharrock & Hughes paper "Ethnography in the workplace" in order to get more information about these perspectives.)

7.1 Grounded theory

One commonly used technique in ethnography is Grounded Theory. Grounded Theory grew out of the 'Chicago School tradition' prominent in the 1920s. Creators of the technique tried to get around the way in which data has often been forced to 'fit' pre-organised categories in more traditional forms of study. As such it's a method for generating categories and concepts from the data itself rather than from a priori categories.

An example of how it might be conducted: you have a topic of study, e.g., how people use Facebook, ask them a bunch of questions (typically open ended), perhaps do other ethnographic collections e.g., participate, observe other people using it. Interviews are often used as the main source of data (which in my experience is the main way it is conducted). Then you try and apply something like the methods below when studying this data set.

There are two main methods here: the constant comparative method and theoretical sampling.

For the *constant comparative method*, we might be studying a topic like nursing care based on interviews with groups of practicing nurses along with some observation we also conduct. When we're looking at our data set, we might develop from it categories like 'professional composure' and 'perception of social loss' which come up in many of the interviews as important topics of discussion. So each nurse might talk about how they perform under stressful situations – they don't have to mention 'professional composure', but we interpret their talk and extract what elementary topic it is exposing.

At some point in the research, the categories also end up developed enough to help guide the researchers as to what to look for, something that cannot be done at the start of the research. So we might start looking for 'professional composure' within later interviews as they are examined for the first time.

A key point is that analysis should continue only so far as it helps develop new categories or properties.

For the *theoretical sampling method*, cases are chosen to provide examples and demonstrations of the theory that is being developed. Cases that are relevant to the theory early on are emphasised, best exemplifying what the theory states; contradicting cases shouldn't be ignored, however.

By either process, a conceptual scheme is developed which becomes progressively more abstract. This develops into what's known as substantive theories (i.e., derived from the data). These substantive theories must be developed whereas formal theories are selected according to their relevance (again, rather than applied a priori).

Thus grounded theory attempts to offer one way to systematically build generalisations from actual data, and has been used in a variety of CSCW studies, and in fact is often the 'default' technique that is applied implicitly within many studies.

7.2 Distributed cognition

Distributed cognition is also sometime used as a basis for ethnographic study. This developed from an influential study by Ed Hutchins on the work of running a ship. Hutchins is a cognitive psychologist, however cognitive science, although diverse, principally is quite individualistic. Generally cognitive science attempts to model the information we receive and the mental processes we go through in order to act in the world, perhaps using theories of representations in the brain and so on. An example might be how we navigate—e.g., comparing mental representations / abstractions with the real world.

This may be the case but it ignores that navigation is often done socially. As a result of studying the cognitive activities of the crew, Hutchins felt that "cognition is a fundamentally cultural process" and couldn't be explained without reference to social groupings. So he extended the cognitive model to consider how cognitive processes aren't contained purely within the individual.

In distributed cognition, cognition and culture are two parts of a larger system; Hutchins felt that cognitive psychology had mistaken the properties of the system for those of the individual. So, understanding the task (in this case running a ship but could easily be a computational activity like CAD or writing software) is a matter of understanding human activity in an environment. This environment contains physical artefacts which themselves are culturally evolved. Distributed cognition's understanding of these environments sees both computational (mental) and social properties.

As Hutchins says "our folk and professional models of cognitive performance do not match when cognition in the world is examined carefully". This brings us to ethnomethodology.

7.3 Ethnomethodology

Ethnomethodology means the study of peoples' methods, and was developed by Harold Garfinkel during the 1960s. Garfinkel felt that sociology was normally done by examining sociological phenomena, e.g., how people order a coffee, using the very methods that members were doing themselves (see later—topic / resource). Social science often has confused topic with resource—it has used members' methods, members' reasoning, etc. as a resource for studying social order rather than making those methods the topic of study. Instead he felt that analysts should be examining the methods that people use rather than layer theory over the top of those methods and call it analysis.

In attempting to solve this problem, ethnomethodology addresses the problem of intersubjective activity. The problem of intersubjectivity is essentially the problem of uncovering how it is that two people are able to, having their own subjective and disjointed perspectives on the world, commonly understand things together. For example, how is it that we are able to converse with one another? How is it that even though our understanding of what words mean has developed potentially in a very separate environment, we are able to make sense of one another? Somehow we are able to bridge our understandings.

Garfinkel, based on earlier work by people such as Schutz, Husserl and Wittgenstein, set about uncovering this. He looked at mundane interaction – i.e., the basic everyday activity which makes up our world.

7.3.1 Things taken-for-granted

Garfinkel recorded conversations between spouses and got people to describe in detail precisely what they meant by each utterance revealing a layer of unspoken-but-assumed common knowledge.

Here is an example he captured of a conversation between a husband and his wife.

Husband: “Dana succeeded in putting a penny in a parking meter today without being picked up.”

(Husband’s description of this sentence: “This afternoon as I was bringing Dana, our four-year-old son, home from the nursery school, he succeeded in reaching high enough to put a penny in a parking meter when we parked in a meter parking zone, whereas before he has always had to be picked up to reach that high.”)

Wife: “Did you take him to the record store?”

(Wife’s description of this sentence: “Since he put a penny in a meter that means that you stopped while he was with you. I know that you stopped at the record store either on the way to get him or on the way back. Was it on the way back, so that he was with you or did you stop there on the way to get him and somewhere else on the way back?”)

Garfinkel also produced a study of an intersexed person called Agnes. Agnes was a biological man living as a woman; her sex was chosen at birth by a doctor because of her having both sets of sexual organs. Agnes felt this was the wrong sex, but had been brought up as male. As a teenager she decided this was wrong and wanted to be female, so began to dress as a woman and take medication and undergo an operation.

This study was an examination of the methods that Agnes used to ‘pass’ as a female. For example, she had to get used to wearing different clothes, acting ‘female’ around others who would have had the assumption that Agnes was indeed female. The way in which Agnes had to achieve gender on a daily basis exposed the taken-for-granted, mundane considerations of ‘doing gender’.

7.4 Core ethnomethodological topics

This section details some of the core ethnomethodological topics.

7.4.1 Breaching experiments

Garfinkel used ‘breaching experiments’ in order to breach the taken-for-granted social order, thus exposing that order for analysis, and thus uncovering how objectivity and normality is achieved (just as Agnes had to use methods for ‘passing’ as a female in order to repair the breach she experienced in changing her gender from male to female).

For example, Garfinkel got his students to treat their family homes like hotels. Obviously this resulted in a variety of incredulous reactions exposing the taken-for-granted aspects of doing ‘living at home’, ‘being a son’ etc. Similarly he got his students to do all sorts of other things in order to breach everyday social order, e.g., staring past someone whilst they’re conversing.

7.4.2 Documentary method of interpretation

The 'documentary method of interpretation' is the treating of an actual appearance as the "document of", "as pointing to", as "standing on behalf of", a presupposed underlying pattern. These "documents" serve to constitute the underlying pattern, but are themselves interpreted on the basis of what is already known about that underlying pattern. Thus 'objectivity' and 'normality' are achieved or constructed via language and bodily interaction, built up through interaction between people. In other words, these are not 'given facts' from nature, i.e., existing, a priori, outside of language.

The documentary method of interpretation is perhaps best described by one of Garfinkel's experiments, in this case it involved a supposed 'experimental' therapy session.

Participants were instructed that they were taking part in a new experimental therapy where the therapist would only answer 'yes' or 'no', whilst sitting behind a screen. The therapist was actually an experimenter who was reading out a random list of 'yes' and 'no' answers in response to the participant / patient's questions. Garfinkel found that participants went to huge lengths to make sense of and interpret the therapists' responses even when they started contradicting themselves. This relates partly to what we saw earlier with Suchman's experience of the 'garden path' problem, as well as relating to indexicality, which is discussed next.

7.4.3 Indexicality

What is this term indexicality? Indexicality is about things dependent for their sense upon the context in which they are embedded. Language and bodily conduct is indexical to a current situation or context. Talk and action have to be related or 'indexed' to current context.

An obvious example would be words such as "you", "us", "here", "that", "next", etc. Bodily conduct that is indexical might be things like pointing or nodding. But there are many more than just these, and indexicality is an encompassing feature of all interaction.

For example, let's think about instructions and procedural guidance (remember the discussion of Plans and Situated Actions, and the procedural guidance of the photocopier's program). Say we were fixing a printer. The system might say to us when it detects a jam, "press the lid button, lift the lid, remove the tray". But these instructions rely heavily on contextual features, and important references can't all be set out in writing in advance. There is a need to index such references, and to adapt with activity.

Consider how indexicality presents problems when designing interactive systems—if you don't include contextual referents then the sense of the activity is lost. This again relates to the problems we saw in Plans and Situated Actions (e.g., garden path problem and the context 'slippage').

7.4.4 Accountability

Accountability is another key topic in ethnomethodology. Accountability is about sense-making, and our actions are 'observable and reportable'—meaning they are made accountable to us by others.

So, we organise and shape our behaviour so that it makes sense to others—we make our activities accountable—extending beyond just the commonsense use of the word 'accountable'.

Being a competent member of society is being able to conduct activity in an acceptably accountable manner—to understand the normative order of activities and conduct yourself correctly according to them. For example, in conversation we

provide opportunities for the taking of turns. This is done by making the end of your talk 'look like' and ending so that the second speaker can start talking.

We can also think about how the breaching experiments exposed this accountability—actions no longer made sense so people involved were offended, etc.

7.4.5 Doing queuing accountably

Accountability is reflexive as well. In order to understand this aspect and accountability in general, we will use an example of a queue.


A physical queue of people exhibits some orderliness (look at the picture) both to members of the queue themselves and those who may observe it (perhaps wanting to join).

Essentially we have a collection of bodies that is 'formatted' into a queue. There is a front, middle and back of the queue. All these aspects are maintained through the orderly observable and reportable activities of the queue's members. Think about what happens when you're queuing—all the ways everyone carefully monitors the state of the queue


Doing queuing accountably means the person joins the queue at the back, rather than ignoring it entirely and heading for the front of the queue. Doing this is conduct that provides an account to others both in and out of the queue of that person's recognition of, and orientation to, the visible, observable-and-reportable orderliness of the queue—so doing queuing also involves saying to everyone else "I am observably adhering to the rules of queuing".

So, this simple action, i.e., you joining the queue, is made accountable, acknowledging to others how the queue is organised through your and their collective action. Reflexivity here means that the very actions involved in doing queuing accountably are the selfsame actions that make up the queue. I.e., accountability is something we do, but it also makes up our very actions – here queuing-accountably *is* queuing.

We can see this when unusual situations expose the assumed, unspoken-but-adhered-to orderliness of queuing. This exceptional queuing situation is recounted by Eric Livingston, an ethnomethodologist, and was something he experienced—maybe you have too? There are three service counters with people queuing up, let's say it's a post office. What happens here is the following:



1. The counter on the left opens up as free.
2. Someone walks over but needs help carrying the boxes to the counter so the person behind places their boxes on the ground and helps carry the other person's boxes to the counter.
3. While this was going on the leftmost counter opens up, the service person starts looking around for the next person in line and the queue begins to bunch up (see bottom right).
4. The queue becomes a 'questioning queue'; people are wondering why the next person in line is not going to the service desk—the orderliness of the queue has broken down.
5. But then the person on the right helping the other carry boxes returns to the head of the queue and picks up their own boxes and heads to the open counter—the queue returns to 'normal'.



The point is that the activities of queue members are accountable to one another; the person doing the help carrying the boxes momentarily broke this normative, accountable order, and the questioning queue was produced. The main thing here is that members of the queue were able to make available their concerns about who was next simply by moving slightly—their actions were then made accountable to the person helping.

We can think about examples from technology which compare and don't compare. For example, emails leave an 'audit trail', so actions often take place in email so that it is accountable. However computers are often frustrating when they are not accountable, such as when they do not give any idea of what they are doing e.g. "please wait"—please wait for what precisely? This is also the case when computers behave in unexpected/inappropriate ways which are opaque to the user. We are in Plans and Situated Actions territory once again, in which Suchman's photocopiers were not accountable resulting in 'garden path' and 'false alarm' problems.

7.5 Other ethnomethodological policies

There are a number of 'policies' which ethnomethodologists typically adhere to.

The *first time through* policy involves treating the topic of study as strange—assuming no knowledge. This involves asking stupid questions, attempting to avoid inference on the part of the analyst, and, it could be said, approaching the topic of study like being a child—a child comes to know and understand what to do in the world—we have to go through the same process for our topics of study.

Precise description is important in the first time through policy. You have to pick up all the things you miss because you're ignoring backgrounded, taken-for-granted activities.

Then there is the policy of *indifference*. Ethnomethodologists aren't interested in constructing social theory—they are indifferent to the variety of theory that can be layered onto social phenomena. As Garfinkel states: "[Ethnomethodological indifference] is a procedure of not needing to consult the corpus of classic methods and findings with which to carry out the tasks of EM research. For this time being we'll carry out the tasks of our research while abstaining from the use of the classic corpus of findings, policies, methods, and the rest. The policy does not advocate the abandonment of established studies. It is a research practice. [...] [It is] not to decide in advance what the phenomenon consists of on the basis of prior formal analytic studies."

Ethnomethodology also studiously avoids mixing *topic with resource*. This means that what is a resource for other social theory is actually a topic of study for

ethnomethodology. For example, understandings of politics might be used to interpret a set of data, e.g., “the politics of “, looking at political affiliation and which supermarkets you buy at—the myriad political theories (e.g., “a marxist interpretation of shopping”) you can choose to understand this data with. Ethnomethodology’s policy is that these two are mixed up. The topic of study should be these resources that social scientists draw on in doing their work of constructing theory.

Social science has drawn on everyday reasoning in order to construct theory, but has missed understanding how everyday reasoning is actually, really done. So, this is an attempt to get away from accounts that ignore the taken-for-granted, background features of everyday life.

Finally, *unique adequacy* is a fancy term for the importance of the analyst to develop competence in their domain of study.

A lot of these policies fit in with ethnography well – particularly precise description and unique adequacy.

7.6 Video analysis

Video is an important resource within many ethnographic studies (and often non-ethnographic ones) in addition to participant observation, notes, documents gathered, etc. Video is replay-able, and enables us to focus very closely on bodily conduct and talk as it happens and then repeatedly review what did happen. Many things are missed when just doing observation and video gives us this extra resource.

Again there are different styles of videoing; we can consider how they’d be used in ethnography. Hand-held cameras might be used for shadowing a person for a day, following an object through some process, and possibly following an activity. Stationary or ambient camera work might be used for recording a particular setting or possibly some activity.

To make things more concrete we can recall the examples given in the Heath and Luff study on the London Underground control rooms. Below is a couple of stills from the paper (taken at the moment the DIA picks up the announcement phone).



Transcript:

C: Control to the train at Charing Cross South Bound, do you receive?
D: Two Four O Charing Cross South Bound
C: Yeah, Two Four O. We’ve got a little bit of an interval behind you.
Could you take a couple of minutes in the platform for me please?
D: (()) Over
C: Thank you very much Two Four O.

(5.2)

DIA: Hello and good afternoon Ladies an Gentlemen. Bakerloo Line
Information...

From the paper:

The DIA does not wait until the completion of the Controller's call before preparing to take action. Indeed, in many cases, it is critical that announcements are delivered to passengers as Controllers are making adjustments to the service. In the case at hand, as the call is initiated, we find the DIA progressively monitoring its production and assessing the implications of the Controller's request for his own conduct. The technology, and in particular the fixed line diagram, provides resources through which the DIA can make sense of the Controller's actions and draw the necessary inferences. At the onset of the call he scans the fixed line diagram to search for an explanation, or provide an account for, why the Controller is contacting a driver and potentially intervening in the running of the service. By the Controllers second attempt to contact the driver, the DIA is moving into a position at the console where he will be able to reach the operating panel for the Public Address system and if necessary make an announcement. On the word "couple", at which point he can infer the potential delay that passengers might incur, he grabs the microphone and headset in preparation for the announcement. In consequence, even before the Controller's call to the driver is brought to completion, the DIA has set the Public Address system to speak to the passengers on a particular platform and is ready to deliver the announcement.

The point is that it would be hard to capture this without the replayability of video.

There are of course various relevant concerns with how video data is gathered and then used. Setting up cameras does impact peoples' interactions—you have to get them comfortable and adapt the filming style to the setting. Also, cameras don't capture everything—video data should be interpreted with this in mind; it's easy to fall into the 'video is all there is' trap.

8 Ubiquitous computing

Ubiquitous computing is a related field to CSCW. Inherently it is about everyday interaction, so the insights of CSCW and CSCW type evaluation are valuable. This will just be a general overview.

Topics that are important in ubicomp include context, everyday interaction, ambiguity (although can apply to all interaction), appropriation and adaptation (like we saw before), seamfulness and seamful design.

But overall ubicomp is about computation embedded in the world. It's about developing tools we can use in everyday 'transparent' ways. Because these tools are to be embedded in everyday life we need to design things that permit us even more to act through them, to help them become unremarkable in their use

Ubicomp's concerns are variously:

- Embedding computational media in 'older media': streets, buildings, furniture, desktops, books, clothes.
- Embedding them into the ways that we work, act and live: part of learning, playing, family, friends, the home.
- Representing more of what people do—input and output is bigger than just input and output.

- Rich input to the computer: multimodal interaction (which we've seen already), but also taking into account our surroundings, our context—so contextual input too (whatever that means).
- Rich output to us: again, multimodal, but also output into the environment of ubicomp devices around us.
- Rich ways for people to present themselves to each other, mediate via devices.

This all means richer human computer interaction and human human interaction


8.1 The original vision

This is how Mark Weiser first talked about what ubicomp is:


The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it. More than 50 million personal computers have been sold, and nonetheless the computer remains largely in a world of its own. It is approachable only through complex jargon that has nothing to do with the tasks for which which people actually use computers. My colleagues and I at PARC think that the idea of a “personal” computer itself is misplaced, and that the vision of laptop machines, dynabooks and “knowledge navigators” is only a transitional step toward achieving the real potential of information technology. Such machines cannot truly make computing an integral, invisible part of the way people live their lives. Therefore we are trying to conceive a new way of thinking about computers in the world, one that takes into account the natural human environment and allows the computers themselves to vanish into the background.

(Mark Weiser, The Computer for the 21st Century, Scientific American, September 1991.)

Weiser helped pioneer original ubicomp vision. How does that vision compares with now? Weiser and Xerox PARC foresaw many recent hardware developments over 15 years ago. Xerox PARC set up a programme to investigate computation by the inch, foot, yard. The rationale was that to support real world tasks we need devices (computers and sensors) of all sizes.




Computation by the inch was investigated in terms of wearable and embedded sensors (see images above), and tabs like PDAs. Compare this with devices like the iPhone, or the Nike+ SportBand. The original vision imagined us having many of these devices with / around / on us.



The next size up was pads (see above)—computation by the foot. PARC developed basic touchscreen computers, and readers. Now consider existing document readers and tablet type PCs. PARC imagined us having several of these.

The final size was slabs (i.e., boards, see below)—computation by the yard. This included things like shared displays, and touch-screens. Now consider modern-day magic whiteboards, the Microsoft surface computing project etc.



8.1.1 Invisible versus unremarkable

Weiser's talk about 'invisible' technology has often been interpreted quite literally, i.e., perceptually invisible—technology embedded into our environment (e.g., in doors, walls, desks, clothes, etc.) so that we literally can't see it. This has to be clarified: it is more that technology becomes unremarkable, becoming invisible in use as part of a daily routine. Consider the unremarkableness of mobile phones: mobile phones are prominent objects, possibly fashionable statements and definitely advertised to us as such, but their use within our society has become unremarkable. It's not that they have just got smaller, they have been appropriated socially.

This notion of invisibility has some roots in philosophy, in this case the philosophy of Martin Heidegger. Heidegger was interested in our ways of being in the world.

One aspect of this is our ways of using tools, e.g., hammers (as he gave examples of) but also we can think about interactive systems like email. Normally tools are ready-to-hand, so his example is using a hammer where it becomes an extension of your arm, and you use it without theorising about it—you forget about the hammer and act through it.

This is linked to psychological notions of 'flow' – essentially another word for being 'in the zone'—think sports, physical activities for most obvious examples. But we can also think about email—we forget about all the technicalities of email and just act through it—most of the time.

When tools go wrong, behave in unusual ways or break they become present-at-hand. So for the hammer the head falls off and needs fixing—we can no longer just act through it as an extension of ourselves—we have to attend to it. Consider the same for email. Everything's fine and we forget about its technicalities until it breaks.

8.2 Ubicomp design

Ubicomp design strives for this unremarkableness. In order to do this it attempts to interweave various digital media into everyday life.

Digital media would be useless and meaningless if we could not relate them to our verbal, written, gestural and spatial activity. As such we must attempt to interweave existing technologies into our everyday life. Ubicomp attempts to improve that interweaving which may exist already to some extent.

Following on from this though, is the question: is interweaving old and new media really so radical? We already do things like reading an email while talking to a friend about it, or, when at a museum we might see a sculpture, read a printed label and listen on a digital audioguide. Or is interweaving just good design? In any case a good system design should fit in its context of use, and fit so well that we don't need to consciously manage, comment on or even notice that fit, even as we use it.

Thus, we return to ubicomp's design ideal: transparency or invisibility. It's not about whether you can see the technology, but whether you have to. As Weiser put it, "so embedded, so fitting, so natural, that we use it without even thinking about it".

A useful and significant ubicomp system is one that we can use in combination with other media. This means traditional media (obviously) but new media too. In each we can refer to and respond to activity in others. What this means is that 'old' media might involve the technology responding to where you are physically, who's with you, what buildings are about, which displays, what furniture etc. that you use. This also means it's responsive to 'new' media as well, e.g., what webpages you've been to, what emails you've sent, digital maps you've accessed, digital documents etc.

A final point of concern for ubicomp design is exposing and sharing activity. This might mean making your past activity a resource for your current activity, or making other people's activity a resource for your current activity. There are many ways to make some activity like this into a resource:

- Make it directly perceptible to yourself and others, so people can use, share, discuss and react to it
- For example, you might make maps, trails, logs available—histories of people you've met, journeys you've taken and so on

You can also analyse and filter this information automatically, so use it as a resource in order to recommend people, places, objects, URLs, and so on. See the Mack Room and George Square examples later.


8.2.1 Weaving computing into everyday life

This example is a game called Yoshi. This was a simple game played on PDAs over the course of several weeks. In the game secured wireless networks are 'yoshis'; yoshis have feeding preferences. Unsecured wireless networks are 'plantations' which yield fruit to feed to yoshis. Players competed for the highest score.

Left image is of a Yoshi whose favourite fruit is melon. On the right is a map showing plantation and Yoshi locations, as well as unexplored areas (question marks).



The graphs below are: (Top) maximum time spent playing the game in one session. (Bottom) daily average of game sessions per player.



In the study players were categorised into two 'modes' of play:

- One method was to change the patterns of everyday life by deliberately setting aside time for special, often relatively prolonged, game sessions, for example during the evening or weekends.
- The second method involved augmenting daily routines by interweaving the game with normal activities, such as work and journeys, and consequently playing for shorter times, but with larger numbers of play sessions.

The winners of the game adapted flexibly to these different modes of play: on workdays they largely augmented their existing activities, whereas at the weekend they changed their normal activities through a specially organised session.

Players reported that the game had some impacts on the patterns of their lives, e.g., work; some gained an advantage by being able to play at work where wifi was available. Some participants would take a different route to their destination, either for work or leisure, in order to play Yoshi. For instance:

- A quotation: “I’d take slightly different street routes than I’d normally take; initially to see what was there. Once I realised there was good stuff there then I’d adjust my route.”
- Another player specifically brought in an orange seed from her home neighbourhood because she remembered that they were difficult to find in her work neighbourhood. After she had seeded an orange plantation, she could feed almost any Yoshi in her work area when she went out to lunch.

Interweaving the game with everyday life also meant managing interactions with non-players including family, partners, colleagues and strangers. This could be positive and negative:

- When visiting family relatives in another city, one player shared the game with them, and went out to play the game with them in their neighbourhood.
- Since access points are sometimes difficult to connect to, he had developed a pattern of walking a couple of metres, then standing still, and then walking again. The player reported how playing the game as he walked with his girlfriend tended to annoy her.

Perhaps most importantly the study found that players built up an understanding of precisely where Yoshis ‘lived’, and where they could be detected. Players considered the distribution of wifi and built up a picture (cf. CYSMN). They figured out areas that were good places to play, and were able to roughly predict whether an area might have many Yoshis or plantations.

In these various ways everyday life, understanding of technology and so on were woven together.

8.3 What makes up context?

A key topic in ubicomp is context, however context presents very significant problems for ubicomp. We can think, for example, about the variety of places we might take a mobile phone: we don’t want it making sound in a theatre, it has to be context sensitive (currently done manually).

Hopefully we have understood that context is not objectifiable. Remember the notion of indexicality—i.e., the importance of context in making sense of indexical talk, indexical gestures and all sorts of other activities.

Context can be thought of a set of features that one can observe ‘now’—but what features? Here is what Schilit, Adams & Want (1994) say about context:

[Context] encompasses more than just the user’s location, because other things of interest are also mobile and changing. Context includes lighting, noise level, network connectivity, communication costs, communication

bandwidth, and even the social situation; e.g., whether you are with your manager or with a co-worker

and Dey et al., 2001:

...typically the location, identity and state of people, groups and computational and physical objects

Let's reflect back to Grudin's interface diagram presented earlier. The diagram has a user, but their interface consists not just of the interface software and I/O devices but also system administrators, documentation, colleagues and so on. Thus, interaction is more like being part of constellations rather than a dyad. Then we can think how ubicomp's understanding of context blows open the doors somewhat and widens this already widening picture.

The main problem with context is that no-one can ever quite find the full and perfect set of 'absolutes' for what context is. Adding ever more sensed 'things' is not enough, nor is making the user control them all by themselves. Once again let's think back to what was mentioned about describing how social interaction is done—we can just keep on going with the 'what I really meant was...' description (as in Garfinkel's spouses' talk).

Next key thing is that context is relative and historical. All the perceptible features of our context give us a vast number of possible actions and vast number of possible interpretations for the perceiver of that context. This is especially with regard to the person, and their companions and acquaintances.

Contextual features of our existence are interpreted on the basis of their overall pattern or configuration relative to each other—i.e., a history of contextual states—making it even harder to do computationally. Interpretation is key. Interpretation may be a learned, rapid reaction or some kind of rationalising, careful thought. For both of these kinds of interpretation, both rapid and careful reactions to contextual features are developed through your historical experience: experience of similar patterns of activity you see, features of the world, people you interact with.

8.4 Ambiguity

Ambiguity is also a central problem we come to when dealing with context. There is great ambiguity in contextual data—e.g., what does it mean to be in this location at this time? Unless we know more then little can be inferred.


But ambiguity does not always have to be problematic, in fact ambiguity can be used when presentation data to users. Ambiguity can be employed to give the user information that reflects the true nature of the information in the system, i.e., it's better to deliver ambiguous information rather than incorrect information or no information at all. This is about exposing the uncertainty in the system to users, being honest about it with the user.

Obviously, ambiguity is not a solution in mission-critical or safety-critical systems.



The images above depict GPS traces of someone heading from the QM into a building a short way down the road. On the left picture we can see information just stopping when the person enters the building – this can be frustrating to users. On the right we can see the GPS trace being incorrect (i.e., we are not suddenly moving north at speed – we are actually inside the building). When information is incorrect it can be confusing. (Remember the adaptation and appropriation of runners in CYSMN—although they started to make sense of incorrect and patchy data they took professional time to do so.)

Alternatively we can use ambiguous information which can often be a more appropriate choice (see image below).



As Bill Gaver states:

By impelling people to interpret situations for themselves, it encourages them to start grappling conceptually with systems and their contexts, and thus to

establish deeper and more personal relations with the meanings offered by those systems.

Gaver et al., 2003

Gaver et al. discussed three types or forms of ambiguity, resulting in a series of design strategies. They have relevance for ubicomp systems where ambiguity is potentially everywhere. We'll look at a couple of design strategies.

8.4.1 Information ambiguity

The first strategy is using ambiguity of information as it is presented to the user. So we compel people to make sense of a system themselves and through this gain a deeper understanding of the system—we get them to do the work (think back to documentary method of interpretation again).

As part of this we use imprecise representations to emphasise uncertainty. Here the system can cast doubt on sources to in order to provoke an independent assessment by the user – i.e., they flag up uncertainty. The system could alternatively over-interpret data to encourage speculation, or expose inconsistencies to create a space of interpretation.

Going back to our GPS example—GPS is inaccurate and often not available. We could try to ignore these problems, always display the position as a point on the map, when GPS is lost just show the last known location. Or, using these strategies, rather than representing position as a set point on the map, draw attention to a region the user may be in—we cast doubt via an imprecise representation of information.

8.4.2 Relationship ambiguity

The second strategy is an ambiguity of relationships—between us, between us and devices, between us via devices.

This allows users to project their own values onto a design. The idea is that they form a more personal relationship with the object or system.. Here we might offer users taking on a role they're unaccustomed with in order to encourage imagination, so in other words changing our relationships to the system—we're not necessarily just a user doing a task. The system might point out things without explaining why—i.e., be more expressive and ambiguous in the way it relates to the world and us. We also might introduce disturbing side effects to create questions over the responsibility of the user.

8.4.3 Ambiguous design example

The History Tablecloth was one of the demonstrations of the ambiguous design technique. The History Tablecloth is designed to cover a kitchen or dining room table. When objects are left on the table, the cloth starts to glow beneath them. This creates a halo that expands slowly. When items are removed, the glow fades. Thus it shows where items have been left over time, making visible relationships between people, and objects, creates questions regarding who has been present and what they have done.

8.4.4 A warning

But let's remember: ambiguity isn't a virtue. It's not an excuse for poor design and not an excuse for creating unaccountable systems. But if ambiguity is applied in a thoughtful manner, it can create more engaging and thought-provoking systems. It provides us with more strategies in our toolbox for ubicomp.

8.5 Seamfulness

The final topic is seams in ubicomp design. A seam is break, gap, or 'loss in translation' in a number of tools or media, which notionally are designed for use together as a uniformly and unproblematically experienced whole. For example wifi networks—wifi is not seamless at all but is designed and implemented assuming that it should be.

Through revealing the seams—the limits and failings in the infrastructure—users can better understand the infrastructure. They can realise where problems are and then ideally they can understand how to workaround or overcome the problems. The main idea is to expose greater resources to users—make the systems more accountable. Remember the discussion we had about ethnomethodology's interest in accountability.

Within ubicomp this idea of 'seamlessness' has been misused—as we've seen people interpreted Weiser's idea of invisibility in the wrong way. Seamlessness in ubicomp systems has been targeted at integration of various components—networks of devices.

This is instead of or as well as interactions between those components. So the net result of this is that we end up making systems where everything is the same rather than letting everything be itself, and work seamlessly with other things. In the end it's easier for us to make seamless systems which use the same components, same software, same design. It's way harder to design systems that integrate well in a seamless way and retain the individuality of those systems.

So Chalmers et al. suggested a "seamful" form of design. The idea here is that seamful design exposes seams as resources for users to draw on. What may be a 'bug' or 'glitch' or even the natural limits of technology can be transformed by systems into an active resource. These resources are made available to users for reasoning about breakdown. This relates to the ideas about ambiguity we've seen.

8.5.1 An example from the everyday world

An example is how mobile phones expose signal strength as a resource. A phone that did not reveal signal strength information would be confusing (not accountable) when you could not make a call—it would have the kind of problematic opacity we talked about before. By revealing this infrastructure information users can better understand and work around the problem.


For example you might purposefully move to a high signal strength area before making an important call. You might even use the seams in technology as resources in conversation to make excuses—e.g., your signal 'cut out' when you were talking to someone you didn't want to talk to. Seams provide plausible deniability in this situation.

8.5.2 Treasure


We can look at an example of a game developed in Glasgow that takes seamfulness as a design technique for the game mechanic.

The game is called Treasure. Players carry a GPS unit and PDA with wifi capabilities. The GPS unit continually tracks the user's position. The PDA continually scans for, and samples, wifi networks and signal strength. The PDA displays the game interface to the user, and the user can use this to pick up coins, upload coins for points, and steal coins from other players. The interface shows wifi coverage and other players' positions (see screenshot below). Players collect 'coins' that are dropped randomly into the world (both in network areas and outside). Players can 'pick-pocket' from one another whilst in network coverage to steal coins. Uploading coins to gain points is not completely reliable, players have to ensure they are in an

area with good network strength to ensure their coins are not lost. Players double their points if they upload coins simultaneously with a team-mate (this encourages collaboration). Team-mates who work together to ensure they upload their coins at the same time gain double points. Players can 'pick-pocket' from opponents by moving close to an opponent while in network coverage and clicking a 'pick-pocket' button.



The game could be played at any location where there was wifi infrastructure networks. Below is a map the players built up of wifi signal strength around the university.



But where are the seams in this game?

The relationship in the game is reversed: any lack of network coverage is no longer a problem, but instead is an integral part of the game, thus there are advantages in being in network coverage and in being outside coverage:

- When in network coverage players learn about new coins, and can upload coins for points.
- When outside network coverage players are safe from pick-pocket attacks, and can continue to collect coins.

Players learned that GPS was not completely accurate. They used this knowledge to both hide from others whilst in network range (e.g., under trees) and to improve positioning. One player even attached their GPS antennae to a metal fence to increase the reception, gain an accurate location, and collect nearby coins.

By employing seamful design, an issue that is normally problematic (not having wireless coverage) can be made into a positive and useful element within the system, even a game mechanic. Thus players become aware of the seams, since simply taking part in the game makes players aware of where wifi coverage in the area can be found.

As coins can be lost if uploaded in areas with weak signal strength, players learn where the strongest signals are. So players learn the underlying characteristics of wifi and GPS, and how they can use these to their advantage

So seams in Treasure open up a bunch of interesting aspects of seams. Developers' normal inclinations are to hide the seams, and attempt to create systems that work "seamlessly" but often this does not work. Problems in the infrastructure do not have to always lead to poor performance, system lockups, or other problems. So by revealing the characteristics of the underlying infrastructure, users will better understand it, and be able to work around problems when they occur. Treasure demonstrated these seams in hardware infrastructure.

8.5.3 Key points of seamful design

- Take advantage of errors, gaps, limits, differences, translations both in hardware and software or data structures.
- Selectively show variation, imprecision, disconnection, categorisation, ownership, cost, past use.
- Accepting that these are essential characteristics of systems' use.
- Design to support users' perception, adaptation, appropriation.

Ideally this helps people find their own way to use systems.

9 References

1. Heath, C. & Luff, P. 1996. Documents and professional practice: "bad" organisational reasons for "good" clinical records.
2. Schmidt, K. & Bannon, L. 1992. Taking CSCW seriously: Supporting articulation work.
3. Grudin, J. 1990. Interface.
4. Heath, C. & Luff, P. 1992. Collaboration and control: Crisis management and multimedia technology in London Underground control rooms.
5. vom Lehn, D. et al. 2007. Engaging constable: revealing art with new technology.
6. Bentley, R. et al. 1992. Ethnographically-informed systems design for air traffic control.

7. Crabtree, A. 2003. Informing the Evaluation of Can You See Me Now? in Rotterdam.
8. Flintham, M. 2001. Where On-Line Meets On-The-Streets: Experiences With Mobile Mixed Reality Games.
9. Sharrock, W. & Hughes, J. 2001. Ethnography in the workplace: Remarks on its theoretical bases.
10. Bell, M. et al. 2006. Interweaving Mobile Games With Everyday Life.
11. Gaver, W. W., Beaver, J., and Benford, S. 2003. Ambiguity as a resource for design.
12. M. Chalmers and A. Galani. 2004. Seamful interweaving: heterogeneity in the theory and design of interactive systems.
13. Schilit, B., Adams, N. and Want, R. 1994. Context-aware computing applications.
14. Dey, A. K., Salber, D. and Abowd, G. D. 2001. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications.
15. Garfinkel, H. Studies in ethnomethodology. Prentice-Hall Press, 1967.
16. Livingston, E. Ethnographies of reason. 2008
17. Buxton, B.: Conversations with Clement Mok and Jakob Nielsen, and with Bill Buxton and Clifford Nass. *interactions*, 7 (1). 46-80.
18. Dourish, P. Where the action is: the foundations of embodied interaction. MIT Press, 2001.
19. Norman, D. The Design of Everyday Things. 1988.
20. Malinowski, B. Argonauts of the Western Pacific. Dutton, 1961.