Chapter 3

Distributed Cognition in Collaborative Systems

> *'Traditionally, human cognition has been seen as existing solely "inside" a person's head, and studies on cognition have by and large disregarded the social, physical, and artifactual surroundings in which cognition takes place' (Salomon, 1993).*

3.1 Overview

This chapter outlines a distributed cognitive framework to enable the examination of work systems within settings. The framework focuses on the representations involved in information processing because access to the representations involved in activity allows analysts to determine the *resources* used in the performance of problem solving, and consequently, design. Artefacts are the physical embodiments of representations, and the media through which representations are operated upon in the world. Understanding the role of representations and the processes involved in transforming them will therefore give an insight into the nature of the resources used to perform collaborative work and design. The chapter discusses how representations, and the processes that are involved in transforming them, are used in cognitive activity that is distributed over collaborative systems. To support this theoretical analysis of collaborative work, a method for collecting data on distributed work systems is described and modified to focus on the representations and processes used.

3.2 Cognition in the world

As the quote at the beginning of the chapter illustrates, the study of behaviour has been dominated by a psychological perspective on the cognitive sciences. The search to uncover the fundamental processes behind behaviour has concentrated on human mental capabilities and attempts to formulate an architecture of cognition (Anderson, 1983; Newell, 1990). In doing so, the sub-disciplines of cognitive science have been sidelined, in particular, anthropology, de-emphasising the roles of context, culture and history (Gardner, 1985), in favour of a stance focused on unsupported mental processes as the main determinant of activity. The development of computer systems to support human performance in the workplace has provided an impetus for re-examining this stance, because in these settings, individual cognitive effort is only of use when integrated with those of others. This has led to the development of an approach to the study of problem solving and collaboration that can account for the situations that such activities take place in and the resources that are available to the actors. The switch of attention towards the resources external to the mind has moved the study of cognition out of the laboratory and into the world (Norman, 1993), and this radical departure from the traditional understanding of cognition is critical to the position taken in this thesis.

Humans have the ability to not only use information in their environments, but also to create tools, or artefacts - man made or modified objects (Cole, 1990). Tools to aid cognition are known as 'cognitive artefacts' (Payne, 1992 - also section 2.2). Cognitive artefacts include external representations of knowledge in the world as memory aids ('knowledge in the world - Norman, 1988), such as checklists and books. They are also used to augment human cognitive abilities, and include devices such as numeric systems (Zhang, 1992), computational devices such as slide rules or calculators, or a combination of both. By performing simple manipulations on cognitive artefacts, humans can logically process information without performing logic operations in their heads (Rumelhart *et al*, 1986a). A fundamental feature of cognitive artefacts is that they do not simply augment existing human capabilities, rather, they transform the task into a different one (Cole, 1990; Norman, 1993), allowing resources to be reallocated into a configuration that better suits the cognitive capabilities of the problem solver.

Cognitive artefacts do not simply support the cognitive processes of individuals (Norman, 1991a), and an example of this is language. Language is a particular form of cognitive artefact (Cole, 1990), that allows humans to spread their cognitive load over a group of people, changing the task from an individual cognitive problem to a distributed problem dispersed over social space.

Expanding the focus of cognitive activity away from the unsupported individual towards a system of tools and groups of people is a far more appropriate unit of analysis if we are to study a real world activity such as collaborative design. Several methods of analysis, discussed in the previous chapter, have been developed to analyse activities involving such multi-tool, multi-participant behaviour. However, none of them are fully appropriate for the study of what is still fundamentally a *cognitive* problem which involves a problem solving approach to be applied by the agents involved. Problem solving involves a system traversing a 'problem space', by moving through various transitory states towards a goal: these problem states are representational in nature, and analysis must therefore focus on these representational states. One method that makes explicit the cognitive paradigm in this enlarged domain of study is distributed cognition.

Analyses with a distributed cognitive framework have been used to examine the cognitive properties of airline cockpits (Hutchins and Klausen, 1990; Hutchins, 1995b, unpublished), the navigation systems of naval vessels (Hutchins, 1988; 1995a), air traffic control operations (Halverson, 1994, 1995), shared CAD systems (Rogers, 1993), shared database systems (Nardi and Miller, 1989), collaboration between programmers (Flor and Hutchins, 1992), and a fishing community (Hazelhurst, 1994), amongst others. This approach to examining the cognitive properties of multiparticipant systems has a great deal of potential for identifying how such systems act as processors of information. To support such activity systems with novel technology, an understanding their information processing requirements and processes is vital in pinpointing where the application of collaboration technology could both benefit work and be implemented without disrupting activity through removing the resources used in co-ordination (Brown and Duguid, 1994; Halverson, 1994). When developing new systems that involve the transformation of work practices, this second point about maintaining the resources used in co-ordination may be as critical as that of proposing novel technologies. Such an understanding allows developers to determine where change should not occur, and where it does, by providing new media that simulate the function of the original co-ordination resources (see also section 2.5.2).

3.3 Cognition as a social phenomenon

3.3.1 Definitions of cognition

There has long been a debate as to what thought involves, a particularly pertinent example of which is Descartes' mind-body separation in the theory of dualism. 'Cognition', as we know it today, is a more recent innovation, achieving prominence in the 1950's and 1960's (Gardner, 1985), and is separated from the much harder to quantify conceptions of mind, involving consciousness, qualia and affect. Initially, cognition was assumed to be a mental activity, and its earliest proponents such as Neisser (1967) and Simon and Newell (1972) wrote about it as involving single individuals. Neisser (ibid.) defined cognition as

operations
required for a given task) are abstract representations, and as such are not restricted to a single individual locus. Indeed, Neisser's above definition of cognition does not delineate who, or what architecture, the cognition should be implemented in. Problem solving does not therefore have to be performed by an unaided individual: any unit performing these activities could be described as a cognitive entity. Individuals can use elements of their environment in cognition, but perhaps most powerfully, groups of people, using artefacts in their environment could be described in terms of the cognitive paradigm.

Work is not normally performed unaided and alone, and the social aspect to problem solving is recognised by Sproull and Kiesler (1991), who argue that 'the fundamental unit of work in the modern organisation is the group, not the individual', and that many of the important aspects of work are 'organised in departments, sub-units, committees, task forces and panels' (p.25). Problem solving behaviour in work activities must involve a unit greater than that of the individual, who becomes a component of the group's problem solving resources. To study a smaller unit of work than the group will miss many important features of the work where problem solving is distributed over a network of individuals co-operating with one another to achieve a solution. Whilst processing of the information available to the group is analogous to an individual's cognitive capabilities, the architecture of this activity differs because of the different representational properties of the resources available. Here, the knowledge base built up by the psychological sciences is less useful in the analysis of problem solving performance for the extended unit of cognition.

3.3.2 Cognition, representation and communication

Distributing work across a group of agents must involve the organisation of that group to co-ordinate activity through some form of communication. In his study of navigation, Hutchins (1995a) describes the hierarchical system of naval rank and the roles that these individuals are expected to play in the navigational fix taking cycle. He documents the representations that communications are encoded in, and how the combination of all of the interacting parts of the system operate to process information to achieve the navigational system's goal (locating the ship in two dimensional space). At no stage in the process can a single person be said to be navigating the ship, which occurs as an emergent property of the individual behaviours of the navigation team. Although the process is not controlled by one person, the performance of the system is not entirely random: individuals are assigned responsibilities and perform roles determined through the prior organisation of work.

The ‘systems’ based perspective on activity needs to describe all of the features that are present in the system: people, artefacts and most importantly, the means of organising these into a useful unit. We therefore have a cognitive system[1] that is mediated through the expression of features arising through non-neurological mechanisms - a system of socially distributed cognition (Hutchins, 1995a). Socially distributed cognition describes group activity in the way that individual cognition has traditionally been described - computation realised through the creation, transformation and propagation of representational states (Simon, 1981; Hutchins, 1995a).

Communication occurs through the transmission of representations (or symbols) on cognitive artefacts between agents. Language is an example of this: it enables the encoding of a (mental) representation that can be transmitted between agents through the medium of speech. Commonly accessible, physical cognitive artefacts are also used as a medium for communication within systems, although they may not explicitly be used as communicative devices. Thus, drawings created for individual use may have a communicative function. For the communication medium (the cognitive artefact) to be used in problem solving behaviour by the distributed cognitive system, the representation encoding the information must have a universally understood meaning between the sender and recipient. Universal comprehension of the medium may derive from common experience, training, or through its use in the setting. This meaning (or mapping) may be self evident, mirroring features of the environment, such as a picture (analogue representations - Woods and Roth, 1988), or they may be more abstract and complex, like text, which require transformational rules for interpretation. Precisely what form these representations take will be determined by the situation that the activity is carried out in, because behaviour is dependant on the resources at hand in the setting, as demonstrated by Suchman (1987). Where there is a choice of representational media, one or more of these media will be selected from those available, the choice of which will be dependant on criteria such as past experience with the artefact and appropriateness to the situation.

[1] It is a cognitive system because it exhibits ‘intelligent’, purposeful behaviours in problem solving and information processing.

Language is an important resource for communication, either spoken or written[2]; however, it is not always the most appropriate medium of representation. In some instances, other representational forms may be more appropriate for communicating information through the system, examples of which may include charts, graphs or drawings. Language is used in combination with other representational media, which provide an indexical focus (something to look and point at) in conversation. A physical artefact can also provide an enduring record of the communicative event to refer back to at a later time, a feature that speech fails to capture. Language is therefore not the only representation carrying communication mechanism that should be examined in the study of communication. Other representational media are equally important in examining communications within socially mediated cognitive systems, particularly ones that operate in media rich environments, and where agents within the system are widely distributed over several distant locations.

3.4 Distributing Cognition

3.4.1 Rationale and aims of distributed cognition

Distributed Cognition (DC) provides a means of describing how the structure of the world, embodied in artefacts and the situational context, imposes *constraints* on the behaviour of the extended cognitive system (comprising of multiple agents within an informationally rich context). The form of distributed cognition advanced by Hutchins is explicitly aimed towards the re-development of systems through their technological media and internal organisational (Flor and Hutchins, 1992; Rogers and Ellis, 1994) to better take advantage of human capabilities and provide support for their limitations. Its goal is to extract information that system designers require in order to make better informed judgements about the information processing requirements of systems of collaborating agents. DC analyses achieve this through deriving the external symbol system (Newell and Simon, 1981) that captures the elements of processing (representations and processes) that transform system inputs into outputs for particular tasks.

There are a number of variants on distributed cognition, ranging from versions where cognition is used as a *metaphor* for understanding group behaviour (Nickerson 1993; Pea, 1993), to versions where the elements of the work system act as a *physical architecture* for cognitive processing, such as that advocated by Hutchins (1995a,b).

[2] Language is a critical component of communication; however, to do it justice, it would require a great deal more effort than can be provided within this thesis, which limits itself to the propagation of representations across multiple media.

This thesis takes the second approach to distributed cognition. A distributed architecture for cognition allows the analyst to go beyond comparing the cognition of the group to that of the individual, which limits analysis to simple comparisons with the capabilities of individual humans. Examining the group as a computational system allows the analyst to examine the emergent behaviours generated through interactions between its component parts. It provides a unique insight into how technology and the socially generated media of communication act upon and transform representations, and in doing so, perform cognitive information processing activities.

3.4.2 Division of labour

Within the distributed cognitive system, problem solving and design expertise lies not only in the knowledge and skills of the individuals, but in the organisation of those individuals, through the configuration of the tools that they use and their work environment. The cognitive analyses of behaviour within a real world environment therefore no longer requires the examination of an individual's psychological functions, but must examine the larger unit and develop new analytic methods to determine how cognition is distributed socially, spatially, materially, and even over time (Cole and Engeström, 1993).

Whilst the structure of the environment is important in determining action (Suchman, 1987), actors do not passively adapt to existing structures. DC theorists (Hutchins, 1995a; Hazelhurst, 1994) claim that the proactive *structuring* of work activity is central to organising and co-ordinating the actions of collaborating individuals. Groups continuously structure their environments through their actions as they perform work. This structuring involves organising and reorganising the physical and cognitive artefacts in the environment, and generating and transforming the social context that the behaviours on these artefacts occur within. These organising structures are retained as representations (either internal memories, or externally as written rules or checklists), or as constraints on behaviour (embodied in the physical artefacts and work environment). The structure of these constraints in the workplace[3] therefore plays a role in determining the architecture for the information processing activities of the functional unit.

The division of labour is a feature of human behaviour that enables our limited resources to spread out and cover an environment too rich in resources to be processed serially by a single actor. Thus, cognitive resources can be considered to be

[3] The architecture is formed from the microstructure of the environment that action occurs within. This in turn develops from macroscopic structures in the historical developments occurring prior to activity.

'shared' amongst several agents (Oatley, 1990; Hutchins, unpublished), which is the principle behind the division of labour. Tasks such as navigation and engineering design are carried out by multiple agents working together, and the division of their labour, in this case, is cognitive labour (Clegg, 1994; Hutchins, 1995a, unpublished). On large collaborative projects, people are often either assigned roles, or come to the project with existing roles (e.g. bankers or safety officers) which means that they have limited expertise across domains. The nature of this specialisation means that most, if not all workers will be illiterate in at least some areas of the collective task. However, they are able to interact productively with other specialists through the division of labour in a particular pattern.

Work activities may also be distributed over technological artefacts, and these too must be considered in determining how work is distributed. The organisation of this socially distributed (cognitive) labour will determine the system's performance on a task. If labour and tools are not organised effectively, the system's task may be performed slowly, incorrectly, or not at all. In shared problem solving, the collaborating agents must organise an effective distribution of labour to bring together their individual expertise to resolve their shared problem, and they must do this by communicating with each other. Understanding how this division of labour operates is central to our understanding of work organisation and working practices (Clegg, 1994). Developing CSCW tools to facilitate the co-ordination of collaborative activity in design will require the explication of the social processes that lie behind this division of labour so that they do not disrupt existing patterns of labour detrimentally.

Socially distributed cognitive labour will include activities such as planning, information gathering and processing, co-ordination activities and group cohesive maintenance, problem solving and decision making. This division of work over people and artefacts is known as articulation work (Strauss *et al*, 1985). As problems and situations evolve during performance of the task, articulation work must involve an ongoing division of labour (Randall and Rouncefield, 1995) that develops and changes to adapt to the situation. The tasks taken on by the agents may depend on a large number of factors including experience, skill, knowledge, training, location or occupation, and task allocation may be imposed, requested, assumed by or delegated to participants (Strauss, 1985). According to Strauss, the organisation of these elements is maintained by feedback in the form of reporting, where actions are monitored, evaluated and revised.

Despite the benefits in distributing labour across a number of agents, there are a number of costs associated with it. Effort and other resources must be put into organising the units that would otherwise have been available for the task, so there is a natural tendency to organise the systems to perform tasks with resource-inexpensive co-ordination activities. A shared understanding of the task and state of activity on the task is important in co-ordinating the division of labour - the greater the discrepancy in these shared processes, the greater the requirement for (more costly) explicit inter-agent co-ordination. In loosely organised systems, agents must be attentive to the work of others to organise the flow of work (Randall and Rouncefield, 1995) and to co-ordinate their collaborative activities. The physical layout of the task environment defines the distribution of access to information, and is therefore a major determinant in co-ordinating the ongoing division of labour (Hutchins, 1995a).

3.4.3 Inside the cognitive system

Cognitive science provides a useful frame of reference to examine intelligence, problem solving and other areas that are considered to form the basis of human intellect through examining the processes that organise human behaviour (Newell and Simon, 1972; Gardner, 1985; Hutchins, 1995b). Its does so through examining how information is represented within the cognitive system, and how these representations are transformed, combined and propagated through the system (Simon, 1981). The added benefit of examining cognition within systems larger than the brain is that many of the representations are directly visible and do not require the indirect methods of examination that experimental psychology has to use. In essence, the analyst can physically enter the cognitive system (Hutchins, 1995a) to see first-hand the representational activity within that system. However, some representations are invisible to examination, because they are located within the mental domain. In the case of distributed cognition, the level of granularity in the analysis is only concerned with the inputs and outputs to agents, and not their internal representations. This is true for the present study, which examines the co-ordination of work *between* agents.

The development of distributed cognition has drawn inspiration from the PDP - parallel distributed processing (Rumelhart and McClelland, 1986b), or connectionist approach to individual, neuronally based cognition, where the whole pattern of agent activation is the meaningful unit of analysis in cognitive behaviour, and the cognitive system is multiply connected and controlled. Important factors in the processing of information by the PDP system are the constraints of the task as well as that of the processor: there is no distinction between the information being processed and the information processing structures (Norman, 1986).

PDP systems are adaptive, configuring themselves to incoming data, and ‘settling into solutions (Rumelhart *et al*, 1986b); in doing so, such a system ‘exhibits intelligence and logic, yet...nowhere has explicit rules of intelligence or logic’

(Norman, 1986, p.537). This has close parallels to groups of individuals collaborating in a task. The difference that the DC system has to PDP, is that instead of using electrical or electrochemical interfaces between the processing units, it operates through socially mediated protocols between the units of the information processing system. Whilst systems of individuals are not as easily specified or homogenous as PDP systems, the PDP approach does show that self organising systems of information processors can work together to produce apparently intelligent and cognitively functioning systems. The distributed processing approach of both DC and PDP therefore entails a major rethinking of cognition, in which the intimate relationship between the psychological and social phenomena is a major feature (Norman, 1986b). As with the PDP systems, investigation of the (social) protocols that maintain and co-ordinate the individual processors is important in specifying the structure of the cognitive processor.

DC, as developed by Hutchins, has adapted the framework of individual cognition to explain how cognitive resources are organised *within a context*, drawing on actors and other features in the environment to perform problem solving. It is concerned with representational states and the informational flows around the media carrying these representations. The DC framework allows researchers to consider all of the factors relevant to the task, bringing together the people, the problem, and tools used into a single unit of analysis. This makes it a suitable candidate for developing an understanding of how representations act as intermediaries in the dynamically evolving and collaborative processes of design.

Hutchins' framework can be developed further, returning it even closer to its roots in cognitive science. Cognitive science allows system descriptions in terms of the functional attributes of a cognitive processor. By disregarding the specifics of implementation, and looking to the higher level terms of what the system *does* rather than *is* (i.e. a functional description), the most basic constituents of a cognitive system can be said to consist of a sensory system, a system memory, a processor and a means of acting on that processed information if necessary. The are displayed diagramatically in fig. 3.1:

fig. 3.1. Functional description of a cognitive system.

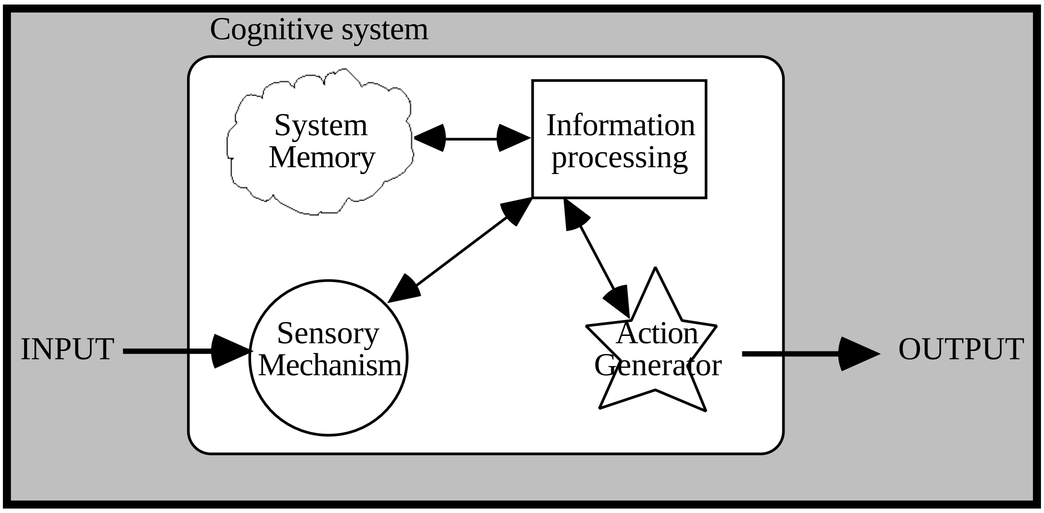

- The *sensory mechanism* takes its inputs (in the form of representations or observed actions) from outside the cognitive system and passes it to the information processing unit in the form of representation.

- The *action generator* allows the production of outputs from the cognitive system in the form of actions or representations. It may also provide feedback to the information processor about the performance of the actions executed.

- '*Memory*' involves the creation of a representational state that is stored to organise subsequent activities; it receives representations from the information processor, and when required, passes them back to it. This storage function may be systematic, or serendipitous, arising through features in the world that are interpreted to inform the system of its past, current and possible future states.

- The *information processor* receives representations from the sensory system and acts upon them, to transform them, combine them or even destroy them. These representations may be stored in the system 'memory', acted upon to create outputs, or used to prime the sensory system to attend to particular inputs.

The implementation of a distributed cognitive system in a real world example can be highly complex: the four units of the functional cognitive system may not fit neatly into individual units - agents can perform several, if not all of the four functions of the system. For example, an engineer may be involved in performing calculations (information processing), they may act as a repository of knowledge (system memory), they may take incoming specifications as inputs to the system (as a sensory mechanism), and they may generate drawings as outputs (action generation).

Nevertheless, the functional description provided within this framework allows the analysis of problem solving with any size of cognitive unit, because the individual is no longer the central focus of enquiry, although they may make up components of the system. It is useful to describe the activity of the cognitive unit using these four components because this framework allows the analyst to understand the functions of the representational states and the activities observed. This demonstrates *what* the components of the system are involved in, and *how* they relate to others in information processing activity.

3.4.4 The unit of analysis

The unit of analysis for a distributed cognitive system incorporates a number of features, including possibly

??

manifested *at the systems level* and not the individual cognitive level (Hutchins, ibid.).

Systems may be interacting, interrelated, or have independent components that combine to perform specific purposes. For example, a car is a system of parts whose purpose is to transport people on roads. It is impossible to understand a system by analysing its individual parts; only through examining the relationships between the components parts is it possible to understand the system. Thus, for a car, an inventory would show that the car was intact even if the spark plugs were in the back seat; to understand why it failed to start, the relationship of spark plugs to engine is important. However, at another level, this understanding of a car is useless, for example when trying to analyse the properties of a road transport system: the grain of analysis is all critical. The emergent properties arising through relationships between the elements of the unit of analysis (in this case, the car) is central to the understanding of DC - ‘the distributed system of cognitions is more than the sum of its components; thus, its operations cannot be understood by examination of its isolated parts, and the system should be examined as a whole’ (Salomon, 1993, p120). The emergent properties of the system arise out of the interactions of parts to generate a new phenomenon larger than the activities of those parts.

[4] The terms functional system (Hutchins, 1995a), functional unit (Rogers & Ellis, 1994) and complex cognitive system (Flor and Hutchins, 1992) have all been used to describe the system under examination, and are interchangeable. This unit of analysis appears to derive from activity theory (Luria, 1979). The functional system is akin to the ‘boundaries’ of a cybernetic system (Rzevski, 1981).

The functional system has emergent properties that cannot be specified by adhering to an individualistic perspective: ‘the unit of analysis is flexible and allows an entire system of actors and artifacts to be considered an intelligent gestalt with properties both similar and radically different from the cognitive properties of individual actors’ (Flor, 1997). The system has different properties to the individuals participating in the activity or system; it, and not the individuals, performs the task and the functional system must therefore be treated (at least functionally) as an intelligent entity.

The goal of analysis is to describe how ‘the distributed structures, which make up the functional system, are coordinated by analysing the various contributions of the environment in which the work activity takes place, the representational media [....], the interactions of individuals with each other and their interactional use of artefacts’ (Rogers, 1993, p. 297). Here, DC and situated action have much in common, because they both consider behaviour to be co-ordinated, at least in part, through an environment rich in organising resources. DC provides a framework to both conceptualise and analyse complex and socially distributed work activities involving technological artefacts and other tools. It operates by focusing on the interactions and actions that are central to co-ordinating distributed work activities (Rogers, 1993).

3.4.5 The role of representations

Within this thesis, the word ‘representation’ is used to describe the way in which a system stores knowledge about a domain[5]. It is a symbolic notion, denoting the thing represent*ed* (Norman, 1991). Representations may be encoded internally in the individual (i.e. mentally), or in the environment (in an artefact). They may encode knowledge about things, or about the organisation of things. They can encode the same knowledge in different ways that are functionally and logically equivalent (problem isomorphs), yet can be manipulated by individuals or systems in different ways.

The importance of the representation in distributed cognition comes from the information processing metaphor of cognitive science, where information is acted upon and transformed computationally (Newell and Simon, 1972). Mental cognition is assumed to be an instance of a Turing Machine, operating through computational mechanisms (Pylyshyn, 1984). Within this computational view of cognitive science, changes to the form of the representations are a part of problem solving, because changing the representation of a problem changes the problem itself. Successive transformations on a representation can transform the initial state into the desired

[5] The word ‘representation’ has a confusing meaning: it can also be used in social science to mean the way that fieldwork is documented and used to ‘represent’ the narrative.

state (Simon, 1981). The computational transformation of a problem state (a representation of the problem) through a ‘problem space’ (composed of the start state, goal state, resources and constraints) into a goal state occurs through the propagation of that representation across various representational structures (Simon, 1981; Kahney, 1993; Hutchins and Klausen, 1991; Hutchins, 1995a,b). In humans, these representational structures would be neural pathways. In DC, cognition takes the form of a computation on a problem representation, involving ‘the propagation of representational state across a variety of [representational] media’ (Hutchins, 1995a, p. xvi; also Hutchins and Klausen, 1991; Hutchins, 1995b), the difference to individual cognition being that the media hypothesised are not limited to internal, mentally held representational states.

Different forms of representational media have particular properties that constrain the uses to which their representations can be put and how they can be accessed. Changes in the medium of this represented information may alter the cognitive state of the system (Hutchins and Klausen, 1991). The forms of representation used, the organisation of the representations and their media, and the interactions of the actors are therefore critical to the task operations performed by the functional system. Hutchins (1995a) provides an example of computational activities and representational transformations in the ‘fix cycle’ of a navigational system, where the navigational system captures knowledge in the world as representations, and successively re-represents them until they can be applied to a chart to represent a physical location (the system goal). The representational states are propagated across a complex set of media, and the goal of spatial orientation is achieved through bringing the representational states of these media into *co-ordination*[6] with one another. Through bringing representations in the system into co-ordination with each other, a representation can be propagated through the distributed cognitive system, being continually modified and processed by a number of individuals and artefacts, until the desired result is reached. Navigation is therefore an emergent property arising out of the combined efforts of collaborating individuals, none of whom can be said to individually determine the course of the process.

Cognitive science has traditionally studied the transmission of representations through the cognitive system using the computational metaphor of information processing theory. Through expanding cognition across a unit greater than the individual, to computational accounts of group cognition, the identification of representational states used by a system can allow researchers to examine the

[6] Bringing the representations into co-ordination with one other involves a process of mapping a representation from one media onto another.

information processing capabilities of the larger cognitive unit. This computational approach is applicable to all areas of collaborative human activity, and will be applied to describe the area of engineering design in later chapters.

3.5 Research methodology

3.5.1 Methodological issues

The analytic framework developed earlier in the chapter provides a basis for understanding how collaborative engineering design work is co-ordinated and through which design work is performed. DC allows the identification of important features operating in collaboration. It also demonstrates how representations are used in information processing activities by the designers in the performance of their work. However, to make claims about how collaboration is maintained and then to fill out the framework with the representations and processes of design work, data will have to be collected about designers involved in real world collaborative design.

As a framework for describing and explaining group cognition, distributed cognition is not a method - its practitioners are therefore able to be eclectic in the range of approaches that they can use. Most of the studies in the literature are observational, although they have been applied in various ways. The analyst is therefore free to select the method of data collection that is most appropriate to the functional system under examination. One such method, ethnography, is applied as the means of data collection in this thesis. However, the ethnographic method was developed independently of the distributed cognitive framework and must be adapted to fulfil its requirements as a tool for data collection. The following sections outline the method used and develop it with respect to the problem domain. It examines the reasons behind the selection of the method, and highlights and discusses issues arising from its use. This is developed in the context of cognitive science and its requirements for the collection of data that is explanitorially adequate.

3.5.2 Research methodologies and cognitive science

The psychological sciences have appropriated the experimental method to examine behaviour, a feature of psychology known by qualitative researchers as ‘physics envy’. Psychology generally uses experiments to answer its questions, but it nevertheless does not argue that reliable knowledge can *only* be obtained through experimentation (Bower and Clapper, 1989). Science methodologies themselves are not solely experimental: astronomy, possibly the parent of scientific method itself, has had to content itself with an observational method (ibid.).

Whilst mainstream cognitive theories have generally taken an experimental approach, in everyday activity, cognition is situated within a social and physical context, and it rarely, if ever occurs in a situation where context does not play a part in it (Cole, 1977; Butterworth, 1992). This accounts for many of the problems faced in designing psychological experiments, where the high number of possible variables must be controlled, and where the environmental conditions cannot be allowed to play a part in the result. However, where cognition occurs *through* the interactions of people and their contexts (Suchman, 1987; Lave, 1988; Norman, 1993; Hutchins, 1995a), as described through the framework of distributed cognition, it is the real world conditions themselves that are the point of departure for the enquiry, and as such, an experimental approach is not applicable.

One approach to examining real world settings in social science is naturalistic research, the approach adopted in this thesis. Naturalistic research allows us to ‘describe what happens in the setting, how the people involved see their own actions and those of others, and the contexts in which the action takes place’ (Hammersley and Atkinson, 1995, p. 6). Indeed, there has been a long history of naturalistic observation in social psychology, because it is good at providing descriptive generalisations about classes of phenomena (Bower and Clapper, 1989).

A cognitive framework is developed below that can be applied to form a basis around which the naturalistic data collected will be analysed. This description begins by describing activities in the terms of cognitive science, framing this in terms of the representations that make up the process. In turn, this guides the methods of data collection, specifying the material required to satisfy the needs of an adequate explanation of the observed behaviours at a cognitive level.

3.5.3 Developing a research methodology for distributed cognition

DC studies the way that work is socially distributed across the functional system, and situated in the context of a physical environment. However, unlike traditional cognitive science, the scientific rigour of the experimental approach (e.g. Schönpflug, 1988) is not possible, nor indeed is it appropriate for the study of collaborative design within the distributed cognitive framework. Appropriate methods of examining the cognitive characteristics of the unit of enquiry must be used, in this case, the unit of enquiry being the functional system. The method of data collection used must be sensitive to the context of activity, so that the interactions within the functional system are accessible to the analyst. However, to demonstrate it relevance, such an approach must clearly define its terms, methods, boundaries and limitations.

The method chosen in this thesis draws from the qualitative methods of social science, yet retains the analytic framework supplied by psychology and cognitive science. This means that the research retains the strength drawn from its cognitive roots, but can go beyond the limitations of mainstream experimental methodologies in cognitive research. This is a new form of research that does not conform to the norms of either psychology or sociology and anthropology. However, the eclectic nature of the approach is not wholly novel; it adapts the naturalistic traditions and methods of data collection, and applies this to a cognitive method of analysis. Only if both of these features can be mutually satisfied can the approach be considered as methodologically sound.

3.5.4 Levels of description in information processing activity

Marr (1982) describes three levels at which information processing systems need to be described to account for a satisfactory explanation of the task. Marr applied these three levels to understanding the underlying cognitive and computational basis of vision, although they can also be applied to distributed cognitive systems (Hutchins, 1995a), such as navigation teams or engineering organisations. The three levels guide the selection of a means of data collection by specifying what is required to satisfy the needs of an adequate explanation of the observed behaviours at a cognitive level. The three levels of description are described in the table below:

Table no 3.1. Marr’s three levels of information processing activity.

Features	**Computational level**	**Representational level**	**Implementational level**
Function	Determines goal of the computation.	Determines how inputs are transformed into outputs.	Determines how computational machinery is embodied in a setting.
Form of description	Makes explicit the entities operated on.	Specifies media and representations available to operate on.	Fieldwork: cognitive description of the system.
Results of description	Specifies high level constraints on activity.	Description of the representations available to achieve computation.	Organisation of components and mechanisms used in transforming representations.

A theory of collaborative engineering design must specify the three areas clearly to achieve a full description of why, on what, and how the cognitive processes operate within the system. By making each layer explicit, the theory becomes open to objective analysis and the possibility of empirical examination (Marr, 1982).

3.5.5 Data collection and distributed cognition

The functional system is the unit of analysis for DC, forming a collection of individuals, artefacts and their relations to each other in a task. Analysis therefore begins by specifying the units involved in the functional system (the representational level), and then by positing a system goal (Nardi, 1992) for the functional system (the computational level of description). Once these have been determined, the means by which this goal is achieved is examined (the implementational level). This will result in a description of the pathways that information flows through, and the external structures that are created and used prior to the solution of that goal, (said to be the equivalent of examining the systems' 'mental state' - Flor and Hutchins, 1992). Information and knowledge can be held as (individual) mental, social and external (including technological) states and may be transformed between the members of the functional system. The work involved in data collection is observational: DC researchers look for information-representation transitions (Flor and Hutchins, 1992) that result in the co-ordination of activity and computations. These occur through the media of knowledge representation.

Knowledge is propagated around the functional unit through a number of communicative pathways (Rogers and Ellis, 1994): verbal; non-verbal; inter-modal transformations (e.g. verbal to text); and by construction of new representations by mental computation in combination with external representations (e.g. operations on tools). All of these can be observed directly. The organisation of this knowledge is important in determining its use, because it constitutes the system's expertise. This knowledge is distributed across the heads of actors *and* in the organisation of tools and the work environment (Hutchins and Klausen, 1991). Data collection must involve descriptions of how the organisation of this distributed knowledge is used in the performance of the functional systems' goals.

There are four areas that require analysis to get a full picture of the knowledge transitions within the system under examination (Rogers and Ellis, 1994):

- by the way that the work environment structures work practice,
- by changes within the representational media,
- by the interactions of the individuals with each other,
- by the interactions of the individuals with system artefacts.

The form of this analysis involves detailed studies of the workplace to analyse the role of technology and work practice in system behaviour. The method by which this is achieved is through performing a *cognitive ethnography* (Hutchins, 1995a), that is, through fieldwork that places emphasis on the representational and representation transforming characteristics of the functional system under observation.

3.6 An introduction to workplace studies

3.6.1 Workplace studies and distributed cognition

The methodological framework for data collection within the analytic framework of DC must identify the representations and processes operating within the functional system. From the data collection, these representations and processes can be brought together in the analysis to describe how they are combined and mediate the collaborative component of work. The requirements of the analytic framework are therefore important in the selection of a method for data collection, because they determine what data is relevant and needs to be collected from the workplace to develop an adequate explanation for the phenomena observed. A work activity that involves actors performing within a complex, real world setting requires a method of examination allowing the study of those actors within that setting. The methods chosen for data collection under these circumstances are therefore bounded with the requirement that an observational study of design work be performed, because removing actors from the context of their activity will change their relationships to the task, to their social interactions with one another, and to their use of tools and technology (the rationale for naturalistic research).

The method of research used in the examination of design, as already stated earlier, will differ from the experimental approach to data collection because of its reliance on controlling the subject and the setting to a limited number of variables: behaviour in the laboratory is not always the same as behaviour in the world. The experimental approach is also limited in that it seeks to answer a hypothesis. This research does not seek to answer a specific hypothesis, or one that was predetermined prior to the research study. Its remit is to examine how work is performed by collaborating actors with a much broader focus of study. Methods of qualitative research are far more suited to the collection of this kind of data.

To understand the differences between quantitative approaches to research and qualitative techniques, their methods are compared below:

quantitative research - associated with the experimental method, where there is a high degree of confidence in the data collected.

qualitative research - research is focused on the context and integrity of the material. The account is not built directly or only from the data collected. Interpretation and subjectivity are central features of the described account of the phenomena, because the research forms part of a debate and is not a ‘fixed truth’ (Banister *et al*, 1994).

Data collection in the field (fieldwork) or workplace (hence, workplace studies) requires a principled approach to provide useful and substantiated information on activity in the domain of interest. Material must be presented in a way that adequately describes the situation, and in a way that the reader can make their own conclusions about the quality of the analytic inferences made from this material. This will entail bringing something from the setting to the analysis, and this traditionally involves descriptions of observed activities and directly quoted discourse. However, to an extent, the qualitative researcher must be 'taken on trust' to provide a truthful description of activity in the workplace. One qualitative research method in particular has the qualities demanded of a research method for the collection of data that meshes with the demands made by distributed cognition: the ethnographic method.

3.6.2 Ethnography - 'making work visible'

Ethnographic analysis attempts to show *how* work is organised (Hughes, Randall and Shapiro, 1992). It has been used to examine the social organisation of groups and has become a central method of analysing workplace activity in CSCW (Hughes, Randall and Shapiro 1992; Grudin and Grinter, 1995; Rogers, 1992, 1994; Heath and Luff, 1991). Woolgar (1988) describes it as: '...a style of research in which the observer adopts the stance of an anthropologist coming upon the phenomenon for the first time. One takes the perspective of a stranger as a way of highlighting the taken-for-granted practices of the natives under study.' Van Maanen (1979) describes it as being used to 'uncover and explicate the ways in which people in particular work settings come to understand, account for, take action and otherwise manage their day-to-day situation' (p. 540). The method is particularly useful in capturing descriptions of socio-historical and environmentally situated behaviours because of its methodologically unstructured nature, which allows a large degree of adaptability in the field. Ethnography also does not limit itself to the examination of predetermined phenomena, which is important for domains where the phenomena are not yet fully understood.

Many ethnographies take place over a period of many years, in which the fieldworker immerses themselves in the domain of study. However, not all fieldwork needs to be extended over such a long timespan, and this is the case particularly where the study takes place in a domain where the language, patterns of social interaction and other features of work are not totally alien to the ethnographer. In the case of a study of a subculture (rather than a true culture) as is the case with designers, the cultural norms differ only in partial degree to those normally experienced by the fieldworker (Bucciarelli, 1994), who more usually operates in a totally different culture, often using a foreign language. This is reflected in CSCW, where there has been a movement towards 'quick and dirty ethnography' and where the fieldwork is of a briefer duration than this extended period of immersion (Hughes, King, Rodden and Andersen, 1994). Because of this decreased level of immersion in the field, the research described in this thesis takes an approach based on the principles of ethnography, whilst hesitating to go so far as to call itself a true ethnography - it involves 'ethnographically informed fieldwork'.

The ethnographic approach allows the examination of features of work that are not apparent from a more cursory observational examination of work practice. Although people may describe the work that they do in clear terms, they do not necessarily perform it in that way (Suchman, 1987; Woolgar, 1994). These are normative (Hutchins, 1995a), or canonical (Brown and Duguid, 1991) descriptions of work: observed reality is not necessarily the same as that described by its participants. This may occur for a number of reasons, although it is likely that much of this knowledge is not explicitly recognised by its users (Rzevski, 1984) - it is therefore *tacit* (Goguen, 1994). Tacit knowledge is inaccessible, and cannot be articulated by those who use it.

Many of the normative rules within organisations exist in documents describing ‘organisational procedures’. However, it appears that these are used as resources for action (Suchman, 1987, 1990), and not as absolute rules that determine behaviour. These ‘rules’ are therefore an ideal of the work process, but should not be confused with the work as enacted. Descriptions of work processes for CSCW design need to make this distinction clear, because there are two ways to describe work: as the way things are supposed to work and the way that they do work (Grudin, 1994). In practice, normative descriptions of work may be abstract, not detailing the exact mechanisms of action, and leaving these to be locally determined by the situational requirements and the resources at hand. The process of fieldwork attempts to prise apart such normative descriptions of work and practice, as performed by actors in situ. The fieldworker must therefore attempt to understand the activity observed in terms of the way that such activity is understood and practised by its participants rather than through simple descriptions of that work.

An understanding of work cannot be gained from a single individual, because they will not be able to describe all of the processes that go on in design because of the distributed nature of that activity. Participants to an activity may not be aware of parts of the work system that do not intrude on their own work areas, so that the task as a whole cannot be viewed from a single perspective, only from a more global perspective, taking in components from the whole activity system. Through an observation of the work, and the communication that links agents involved in that work, a more realistic picture of the design process can be built up that does not rely on the perspective of a single person or their subjective opinions of work.

3.6.3 The ecological basis of the ethnographic method

Ethnographers propose that ‘things’ (the object of examination) should be studied in their natural state (naturalism) - and this precludes the experimental method. Because behaviour cannot be divorced from the situation, an ‘in situ’ study of behaviour must be conducted to understand the unit under analysis (Winograd and Flores, 1986; Bannon and Bødker, 1991; Hutchins, 1995a). The primary aim is therefore to describe the setting (or ecology), context, and both the actions performed, and the way that participants interpret their own actions. This epistemological gulf between ethnographic and psychological approaches to the study of cognition and work has resulted in a long running intellectual debate between psychologists and ethnographers (see Monk *et al*., 1993). When choosing to use ethnography as a technique, the needs of relevance must be balanced against those of trust in their applicability (the qualitative - quantitative debate). This will involve matching the needs of ecological validity (high for ethnography) and reliability (i.e. can the results be reproduced - low for ethnography, but high for experimental approaches) to the domain of enquiry and the nature of the research problem.

Ethnography is characterised by -

- data[7] gathered from a range of sources (data triangulation)
- an in depth study of one or more situations (source triangulation)
- the study of data in context - good at revealing complexity, rather than stripping it away
- an unstructured approach to data gathering, allowing key issues to emerge through ongoing analysis; it does not involve hypothesis testing.

In a naturalistically observed complex environment, such as a workplace, people draw information from a huge number of sources that cannot be replicated in an experimental or laboratory based environment. In particular, observations emanating from experimental studies ignore situational and organisational features (Anderson and Sharrock, 1993; Norman 1993). The ethnographic approach is therefore the most appropriate means of analysing the cognitive activities performed by the functional system of design, and will be used in gathering material for a DC analysis of design. A ‘cognitive ethnography’ can provide a method for discovering the representations (and mechanisms for their propagation) which are operative within particular activity systems (Hazelhurst, 1996). Using distributed cognition as an analytic framework, the ethnographically informed fieldwork can be used to gather material which can then be examined and documented, picking out the salient points relating to the computational, or cognitive, characteristics of the functional unit.

3.6.4 Analysis of the cognitive ethnographic data

Analysis of the ethnographic data will be multidisciplinary, drawing from cognitive theory, and borrowing from the intellectual heritage of anthropology and sociology. This perspective will influence data collection, placing emphasis on the role of artefacts, and collaboration around these artefacts. Observations will therefore centre

[7] The use of the word ‘data’ is problematic because it implies that there are hard facts within it; rather ethnography is a form of ‘reportage’ (Anderson, 1994). However, data is the term used in Agar (1980) and Van Maanen (1979), and the term will be used bearing this criticism in mind.

around the types of artefacts that are used or created, how they are used, who they are used by, how changes are made to them, and how the organisation structures access to these artefacts.

In a cognitive description of a functional system for engineering design, Marr's description at the level of 'implementation' will involve the main cognitive ethnographic component of the work. This involves descriptions of the design group structure, and how the representations described above are transformed to perform the computational functions of engineering design. The cognitive ethnography will involve moving from first order concepts (the observed data), towards second order concepts (theories about the data) that account for patterns in the first order data (Van Maanen, 1979). The distributed cognitive analysis will be the method used to transform the observed data from the fieldwork into a theory of the patterning observed. The process is iterative, involving data collection, the creation of a 'working theory', followed by repeated sequences of observations and theory matching exercises. The result of this is an account[8] of the distributed cognitive processes that underpin the activity of the observed functional system.

A cognitive system is necessary to carry out intelligent actions, consisting of a control mechanism, a memory and a set of operations to act on an input, and to generate an output. The task of the DC theorist is therefore to examine the data and to identify the set of processes that create, modify, reproduce, or transform representational structures in that system. The goal of the researcher in cognitive ethnography is to provide an account of how the distributed structures of the functional system are co-ordinated (Rogers and Ellis, 1994).

3.7 Conclusion

The chapter examines how cognitive science can be used as a means of examining units of activity greater than that of the individual to include several people, tools and the structure of the environment in problem solving behaviour. It outlines and develops the theoretical basis behind distributed cognition and demonstrates how it can be used to guide data collection and analyse settings. The role of representations and processes is examined, and their relevance and importance to the performance and understanding of collaborative activities is highlighted. A method of data

[8] Although not *the* account (Agar, 1980). Ethnographic reportage is a subjective means of describing a situation, and there can be no single 'correct' description - however, some descriptions of work may be more appropriate under certain conditions.

collection, ethnographically informed fieldwork, is also outlined that can be used to collect material for the analysis.

In principle, the framework of distributed cognition is applicable to the examination of all areas of multiparticipant activity. One such area is that of engineering design, a particular instance of collaborative activity. Studies of work, such as those revealed in the fieldwork and analysed within a theoretical framework can be used to reveal the social organisation of activities, the use of artefacts, and the mechanisms co-ordinating the behaviours of collaborating individuals. The following chapters will attempt to reveal these patterns in a study of collaboration in engineering design.