Process, representation and taskworld.
Distributed cognition and the organisation of information.

Mark Perry
C.i.e.*R*., Department of Information Systems and Computing,
Brunel University,
Mark.Perry@brunel.ac.uk

Abstract

Distributed cognition provides a means of describing the co-ordination of collaborative activity. A single framework is applied to examine the interactions between people, the tools they use, and the environments that their activities are situated within. The resultant analyses show how the system resources are applied to result in problem solving activity. Examples from fieldwork are used to explore these issues. The paper critically evaluates DC, exploring the problems and the benefits that such an approach brings to understanding the organisation of information in its contexts.

Introduction

This paper shows a technique for examining information within its context, drawing together threads from anthropology, sociology, psychology, and organisational theory. As stated in the call for papers, the formalisation of 'things' 'into formation' is a partial, precarious and political activity. On the other hand, when applied appropriately, this process can also be a powerful means of achieving *an* understanding about a setting, albeit with a particular perspective. One such approach is presented here of looking at this idea of 'information' as both the product of an analysis, and its use in an organising capacity.

There is an applied need for providing rich descriptions of action in real world settings both in the development of work-appropriate technology, and to support organisational change. This requirement has lead to, and been informed by, the development of theoretical frameworks to organise this information. Amongst others, these include ethnomethodology, situation theory, actor network theory and activity theory. A relatively new framework that allows researchers to organise task related information about activity and context is distributed cognition (Hutchins, 1995a,b), developed specifically to analyse and provide resources for redesigning systems (Rogers & Ellis, 1994) by examining their information processing characteristics.

Distributed cognition (DC) draws from the information processing metaphor of cognitive science, where systems are considered in terms of their inputs and outputs, and tasks are decomposed into problem spaces. However, DC differs from traditional approaches to cognition by taking a broader approach to the unit of cognition, which considers the 'functional system' as a unit of analysis. The functional system can comprise of multiple individuals, working in a complex environment, and using artefacts in performance of a task. A DC analysis examines both the work and the division of labour required to co-ordinate the activities of the individual agents within a single framework. This allows the artefacts of work to be examined as resources for co-ordinating the collaborative aspects of work (Robinson, 1993, Brown & Duguid, 1994), as well as the means of performing work.

In DC, information is studied through the *representation* that it occurs within - the media or artefact that symbolises the information. This is a more manageable unit of examination than the more slippery term 'information'. The context that work is situated within sets the problem and specifies the resources (and constraints) that are available to the functional system: it determines the *taskworld* within which activity can occur. The activity that transforms the representations within the representational media occurs through the *processes* of work, either as structured, organisationally determined procedures, or social practices. Describing work systems in these terms allows the analyst a more easily specified set of criteria for examining and organising situations of activity than using those of 'activity', 'information' and 'context', none of which can be easily categorised of defined. Through the DC analysis, the representational media, the processes transforming the representations, and structure in the world that the agents use to make sense of their activities are made explicit, giving an insight into both the work, and the organisation of the elements of work (the division of labour).

Rich *enough* descriptions of action

Within social groups, DC can be applied to show how interpersonal processes are used to co-ordinate activity, where cognition is said to be 'socially distributed'. The distributed cognitive approach is used to identify information processing activity in a particular task-system through its inputs and outputs, processes and representations. In practice, this information can be gathered in these types of social and organisational systems through field studies that examine the task, goals, participants, artefacts, the various resources and constraints, transformational activities and co-ordination mechanisms that are used in the performance of work.

The framework of distributed cognition gives an analysis of a task phrased in terms of the people, artefacts, context and their organisation with respect to one another - the representations, processes and taskworld. However, *only* the resources and activities that compose of the information processing component of the task are featured in the analysis. Whilst it is recognised that such descriptions cannot capture the full complexity of the situation, concentrating on these relatively simple elements of work and their interactions can provide a qualitatively useful analysis of work. It is an interpretative form of analysis because the described systems are descriptions from the viewpoint of the analyst; they are necessarily subjective and limited to the range of experience that the analyst has access to. This is a significant feature of the approach, and is one where distributed cognition diverges from cognitive psychology, which tends towards describing a system's information processes as *the* method by which the cognitive system operates, rather than one of many possible means of description.

Pushing the boundaries of information processing theory

The information processing (IP) approach to human cognition assumes the representational theory of mind, as prompted by cognitive theorists such as Pylushyn (1984) and Simon & Newell (1972), and where information is encoded symbolically in representational states within some physical medium. Changes to these states are moderated by various processes that act on them, and successive state changes to the representation result in problem solving. This approach is, however, not exclusive to a description of 'mind', but allows the analyst of any such information processing system to describe the activities involved in goal directed behaviour. In this sense, the IP approach to problem solving is hardware independent.

In situations where more than one person is involved in problem solving activity, or where the problem solving process is externalised onto physical artefacts, the IP metaphor can be used as a means of examining what aspects of problem solving behaviour occur in particular activities and physical locations. In these situations, the representations being passed *between* individuals through the IP system are visible to the analyst, in the artefacts that are created, modified and combined. This means that the analyst is able to 'step inside' the cognitive system (Hutchins, 1995a) in a way that is not possible to do in psychological studies of mental activity (Perry, 1998). In the framework of distributed cognition, cognition is not simply used as a metaphor relating to comparisons with the human mind, with which to describe collaborative systems (Larssen & Christensen, 1993). Rather, it is used as a means of describing the functional architecture of a problem solving unit.

Cognitive systems can be, at least functionally, described as consisting of four elements. These consist of a sensor, and an action generator, which enable the system to take in inputs and to produce outputs. They also have an information processing system which can transform the representations either sensed, or stored within the fourth part of the system, the 'memory'. It is important to emphasise that this is a functional description of a system - the framework does not attempt to describe how these elements are physically embodied in the cognitive system, simply that the system can be usefully be described in these terms. This framework is demonstrated in fig. 1 and its components described in more detail below (from Perry, 1998):

fig. 1. Functional description of a cognitive system.

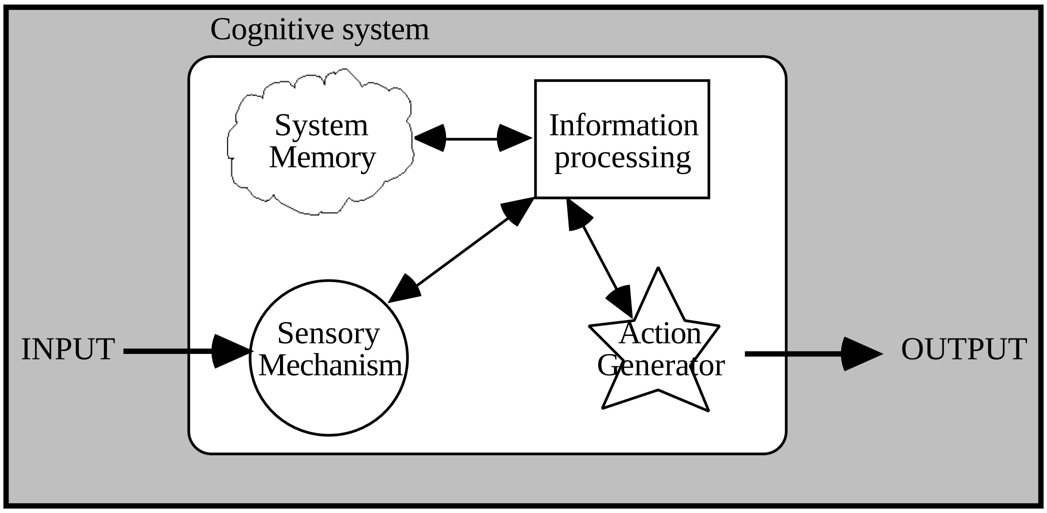

- The *sensory mechanism* takes its inputs (in the form of representations or observed actions) from outside the cognitive system and passes it to the information processing unit in the form of representation.
- The *action generator* allows the production of outputs from the cognitive system in the 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 framework proposes that the IP activity can be described solely in terms of the representations, and the processes that act on these representations. Whilst the four units of cognitive activity described above appear to be relatively simple, we do not wish to trivialise them, or their interactions with one another. 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). This example illustrates the importance of 'granularity' in the analysis - the granularity of the example is given in terms of the individual level of work activity. In other cases, several individuals may interact in a work system, and the granularity will therefore differ. It is often possible to describe complex functional systems in terms of several embedded simpler functional systems.

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.

Representations and processes in work

The distributed approach to information processing provides researchers with a powerful tool with which to explain how co-operative activities are co-ordinated in real world settings. Analysts require a means of describing the components within a social system to explain the mechanisms that co-ordinate groups of co-operating individuals. In cognitive science, these properties are described in terms of the representations and processes of individual thought. This cognitive framework can be expanded to examine larger units, to include individuals interacting with external representations, and the interactions of multiple individuals in a work setting. This cognitive process, as proposed by Hutchins (1995a), involves computations 'through the propagation of representational state across a variety of media' (p.xvi).

The development of distributed cognition has drawn inspiration from the PDP - parallel distributed processing (Rumelhart and McClelland, 1986), or connectionist, approach to individual, neuronally based cognition. In PDP systems, the whole pattern of agent activation is the meaningful unit of analysis, 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). Whilst systems of individuals (as in DC) 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 entails a major rethinking of cognition, in which the intimate relationship between the psychological and social phenomena is a major feature (Norman, 1986). As with the PDP systems, investigation of the (social) protocols that maintain and co-ordinate individual processors is important in the specification of cognitive structure for a DC analysis.

Cognitive systems must be composed of information processing units. In social systems, these agents are invariably people - but people work as individuals - individually functioning, independent agents (Salomon, 1994). However, through co-operation, individuals bring unique skills and resources to problems that they can use in conjunction with one another to solve their shared problems. The co-ordination of these resources is crucial to the co-operative activity that they are undertaking. Each agent brings resources to the problem and must engage in communicating their ideas to the other participants using feedback to modify their behaviour in the light of the other agents' activities. Failure to co-ordinate these mechanisms will result in the failure to produce a good solution, or a solution at all.

Technology and the division of labour

One of the problems in performing collaborative work is organising the task into component parts that can be performed by individuals. This must be managed so that the parts can be physically integrated back together again after the component parts of the task have been processed. Attendant to this process is an issue of *co-ordination* to ensure that the individually assigned parts are performed correctly, and in a form that can be re-integrated with the whole. Hutchins (1995a) describes this as the 'division of labour' and demonstrates how it is mediated through social, cultural and organisationally determined protocols in a navigational context. Difficulties in the process of distributing work can arise through individuals not performing their set roles, but also because the individuals fail to co-ordinate their behaviour to perform the task. Improving the co-ordination of work can be effected either through organisational change or the adoption of new technology, neither of which can be considered in isolation (Grudin, 1993).

Individuals can be aided by developments in technology that enhance their productivity through aiding their creativity, memory, information processing capabilities and other human 'inadequacies'. However, integrating the individual's performance with that of others is crucial to the performance of the group, and this involves co-ordination. New and more powerful communication facilities by themselves will not of themselves necessarily result in better co-ordinated work; however, when used appropriately, they can allow individuals to work together more effectively. To support functional systems with technology, it is also important to understand their information processing requirements so that technology can be implemented without disrupting activity by removing the resources used in co-ordination (Brown and Duguid, 1994; Halverson, 1994). When developing new systems that involve the transformation of work practices, maintaining the resources used in co-ordination may be as critical as that of proposing augmentative technologies.

Field study - Design in civil engineering

To demonstrate how the distributed cognitive analysis can be conducted, an example from a field study is given below. This example is an analysis of engineering design in the construction industry. The complete study is too large to present in its entirety here - it is not an in depth examination of design, but demonstrates how an information processing approach can be applied.

Background - goals and resources

The background of the study is an importance component of the analysis, because this provides the context that problem solving takes place within. The background to the activity is used to identify the functional system's goals and resources. The resources include the people involved, the artefacts available to them, and the organisational relationships between them. These make up the underlying structure of the functional system.

The field study described involved a road building project. Fieldwork covered the involvement of three spatially distant teams working for the construction company (pseudonymously called ConsCo). Several other organisations also took part in the construction process, and these were also examined, although not in such depth. Note how the *work activity* determined the boundaries of the functional system, and not the commercial entities.

To build the 'permanent works', or permanent road structures, ConsCo had to design and build supporting structures known as 'temporary works'. The temporary works comprised of non-permanent items, including scaffolding, concrete moulds, and supply roads. These temporary works originated in the designs of the permanent works, but had to take into account site conditions (such as slope, weather, soil condition, and existing structures) and available resources (money, time, materials, equipment and labour) not accounted for in the initial designs.

A number of people were involved in this design process, including a construction team, made up of a team leader, seven engineers (one senior, three site, and three graduate engineers), two foremen (supervisors), five gangers (junior supervisors), and general labour (varying around forty). The team operated on-site, away from the main site office and were supported by a design co-ordinator. The design co-ordinator worked closely with a design engineer to develop the team's requirements into a temporary works design solution. Several other individuals and groups external to ConsCo were also closely involved, including a 'resident engineer', who checked that the road was built to contractual standards, a railway operator and an environmental agency, each with their own responsibilities.

So, the *goal* of the design system was to construct a road, following the permanent works designs and contractual details, with commercial constraints (cost) without disrupting the railway or causing environmental damage. The functional system therefore included the construction team, the design co-ordinator, the design engineer, the resident engineer, the railway operator organisation, and the environmental agency.

Inputs and outputs

Following from the identification of the goal, the initial inputs to be processed must be identified. Also, the outputs of the system also have to be made clear. The identification of the inputs and outputs helps to identify the steps required to moderate transformations between them. The basic inputs and outputs to the example design system are described below:

The **inputs** included the permanent works drawings, the contractual details for construction, conditions prevailing on the site, paper documentation about the design process (handbooks and manuals), and knowledge about the design and construction process in the heads of the agents involved.

The **outputs** of the design process were the finished temporary works designs, approved by the team, the resident engineer, the railway operating organisation, and the environmental agency.

Whilst these inputs and outputs are at best sketchy, they delimit the structure of the problem that the design system must solve. Once the inputs and outputs have been identified, the processes and representations that are used by the functional system can then be examined.

Representations and processes

A huge number of representations operating in the design system were identified. These were held within the minds of individuals (but not directly observable), and in artefacts, which formed the tangible media representing design information (observable). These media included drawings, paper and pencil sketches, letters, post-it™ notes, schedules, letters, method statements, risk analyses, works records (a diary of site instructions, records and requests for information), and other paper based forms that had to be completed in the course of work. The artefacts carried representations that could be communicated between the collaborating actors involved, allowing them to perform their own individual tasks, as well as the collaborative design goal.

The design processes were partially prescribed in the handbooks and manuals that determined relationships between agents, and in the organisational structures they inhabited. These specified where the responsibilities for tasks lay and determined the roles that the individual agents had to fulfil. Whilst these were often followed, particularly the safety related rules, they were used as general organising principles, followed when appropriate, but worked around when other methods were more appropriate.

The artefacts also provided a resource for structuring information processing activity. In some cases, the artefacts helped to determine how they were to be used, by limiting the functions that could be performed through their physical structure. This structure made certain activities hard to perform on the artefacts, and others easier. For example, paper forms had boxes to fill in or tick, limiting the actions that could be performed on them.

Structures in the environment at the site and office also helped to determine the processes applied to particular problems or representations. These determined the configurations of representations and processes applied to the design problem. For example, on the site, large distances made the location of people difficult, and communications were disrupted by this. In the office, however, because agents were co-located, they had a broader range of communications media at their disposal. The organisation of the construction team's office was central to interactions between the team, because it controlled access to problem solving resources. When information was required by a team member, it could be accessed, either by asking other people in the office, or by searching through the filing archive. Project information was held in files scattered around the office. Individuals also maintained their own files of activities which could provide information to other people when they were not physically present in the office.

Spoken communication from the desks, allowed all of the people in the room to be aware of developments (similar in nature to Heath and Luff, 1991), or contribute to the discussion. When the senior or site engineers wanted to speak to the graduate engineers, they could stand up and talk over the top of the desk partitions, providing a visual and auditory focus of attention in the room. This allowed people to keep aware of ongoing conversations, and abreast of developments. In addition to these 'open' conversations, telephone calls were carried out in loud voices; this was partly because the level of ambient noise in the room could be fairly high, but also because it allowed the others in the room to overhear part of the conversation.

The distributed nature of the site made contacting distant individuals difficult. When people were not present, various media could be used to communicate, through the radio, by placing written notes, sketches, method statements or risk assessments on people's desks, or jotting notes onto a whiteboard. Messages could also be left with other people in the office to pass on when the person returned. When information was required from a person who was not present, paper-based information on people's desks ('desk litter') could provide information to their location or the task that they were currently engaged in by reference to the drawings and other representations on their desktops. Other physical artefacts also provided information on the people's whereabouts: if a person's Wellington boots and hard-hat were not present, they were probably on site; if someone had a pair of wet, muddy boots under their desk, it meant that they had recently been on site and could be asked about it. Even the window was used to see whether people's cars were in the car park; if this was the case, then that person was likely to be available.

Contact between the dispersed team members with the site (distant from the office) was made possible through the use of a portable hand-held radio link, allowing the engineers, gangers and foremen to call each other. The radios were kept on all of the time so that contact calls could be made. Interestingly, the background noise of the radios was also used as a means of indirectly monitoring general activity on the site because distant conversations could be attended to. This was possible because of one of the qualities of the radio as a medium of communication: the radios were set to an open channel, and all communication took place on a common wavelength. This meant that both sides of a communication could be overheard by non-participants with access to a radio. As in an open plan office, which allows overhearing, or 'surreptitious monitoring' of conversations, the radios provided a similar function between different individuals.

Transformational activity

The transformational activity performed by the functional system involved taking the inputs, transforming them into representations that could be applied in the design process, and re-representing them until they matched the desired goal as an output. The transformational activity involved co-ordination of the representations, matching them to an appropriate process. A description from the field study about how these transformations were performed is given below. In this example, information was gathered from the field about the conditions on the site relating to a problem with the bridge.

The example demonstrates how information represented in one rich media is extracted and synthesised into a simpler representation with a structure that is appropriate for a particular task, that of comparing the designs to the physical conditions on site:

A graduate engineer had spent several minutes poring over a drawing, taking measurements of the planned gradient of the surface of the bridge. These measurements were transferred onto a sketch. The measurements on the sketch were in a different format to that of the original drawing: whilst the original drawing had been an overview of the deck (viewed from above), the sketch was a section through the structure (viewed from the side).

In addition to the different aspects of the drawing and sketch, the axes on the sketch were chosen so that they exaggerated the gradient and made deviations and discrepancies in the data more easily visible: the horizontal axis was on a scale of 1:250, whilst the vertical scale was 1:10. Information from the actual conditions on the site, taken with geotechnical equipment and previously recorded on another table, were then also annotated onto the sketch (fig. 1).

fig. 1. Sketch of road gradient.

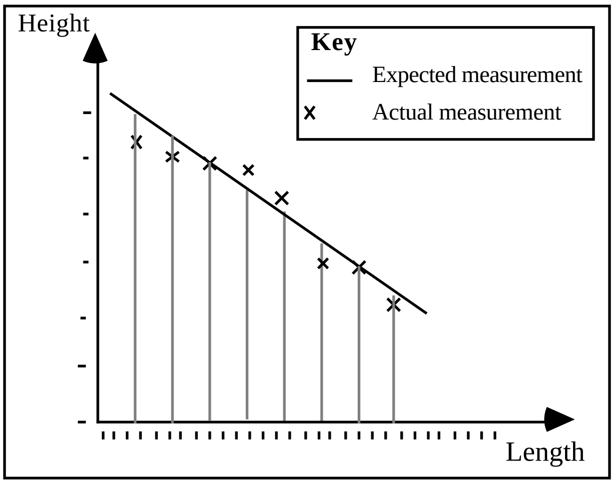

On completion, the graduate engineer left the sketch on the senior engineers desk, with a note attached to it explaining that he had found a discrepancy between the expected and actual gradient. The note further commented that he was going to be away from his desk for the rest of the day, but informed the senior engineers that he would be working at a particular location if further information was required.

The example shows how information was transformed from the design drawings and site onto different representations (tables of measurements), and then onto a graphical representation (the sketch) that more clearly demonstrated the relationships between the two data sets. The sketch shows that the measured slope had a gradient that did not match the gradient on the drawing. The form of this representation clearly demonstrated this discrepancy, through the difference that was exaggerated on the differential scales of the axes.

The reason for this problem was that a sub-contractor had driven the piles to incorrect tolerances. This discovery of had important consequences on subsequent design activities because it limited the loading that could be placed on them. The finding resulted in a change to the construction process and temporary works design in surrounding areas.

Applicability of the approach

The example demonstrates how to perform a distributed cognitive analysis. It shows how a cognitive approach can provide a means of exploring situated activities without the analyst experiencing the 'blooming buzzing confusion' (James, 1890) of the complex observed setting. The approach allows the analyst to break down the observational work into a manageable collection of interrelated features that are relevant to the problem solving activity. By using methods from cognitive science we can give 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), we have shown how the most basic constituents of a distributed 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.

The external symbol system derived in the analysis will be incompletely specified and is *not* intended to be a formal computational structure. In practice, such a method would be too costly to examine for *all* transformational actions in the functional system. In addition, the interpretative nature the ethnographic approach does not lend itself to this approach. The method of data collection and its analytic framework suggested here do however provide a means of examining aspects of information processing by a group. It cannot be treated as a total description of work, but more as a means of getting a deeper understanding about the activity *in its context*.

Whilst DC has a number of highly useful features in representing the (computational) nature of collaborative work, we recognise that, as with any description of activity, potentially valuable information about work can be lost. This can have benefits, when information not relevant to the problem is filtered out, but it also means that important information about the setting can be lost. One example of this is in motivation - DC analyses cannot provide a means of including the motivational factors that operate in work relevant to the performance of the task. The problem solving systems described in DC are 'impersonal' (i.e. systems are described in terms of their external resources), and whether or not the work is tedious or poorly rewarded, this factor is not considered.

In the DC analysis, all agents are expected to perform work in the same way, whatever their personal concerns. This may be a limiting factor in the application of the technique, although it is likely that many of these factors are irrelevant outside of particular work situations. This limitation has an unexpected advantage: when applying the analysis, the technique may be useful precisely *because* it reduces the importance of these setting-unique factors, allowing generalisations to be made. DC simplifies the way that the setting can be conceptualised by reducing the complexity of detail in the setting to that relevant to the information processing element of work. Moreover, as well as structuring the analysis of the field data, the framework also guides the field worker to actively look for particular features in the setting - the features that contribute to the problem solving activity - as either representations, or processes that transform those representations.

Conclusions

DC, as developed by Hutchins, has adapted the framework of individual cognition to show 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 dynamically evolving collaborative processes.

A DC analysis results in a detailed description of the data collected in the workplace studies, with particular reference to the people involved in the work activity, the tools used in performing work, their relationships to one another, the procedures followed, the situations that action occurred within, and social interactions between them. The methods and tools provided by the DC approach allows the analyst to specify the divisions of labour in the distributed cognitive system. The material is structured by attending to the goals, inputs and outputs to the work system, the representations involved, and the processes used to transform these representations. The analysis contributes to a description of both the work and the co-ordination of that work within a single language. The importance of this is that communications and co-ordination are situated within the context of the work itself, and not separate from it. The intention of this is to show how the patterns of organisation and communication generate the cognitive properties of work systems by demonstrating how representations are used, both as a means of *organising* and *undertaking* collaborative tasks.

Analyses of activity using distributed cognition as an organising structure allows the instigators of organisational change to make better informed decisions about the systems they are involved in transforming. It allows systems designers to visualise where problems exist for collaborative problem solving so that the functional systems can be augmented with appropriate technological support. It also allows managers to manipulate the configurations of people and existing functional systems to make better use of the resources available. As importantly as suggesting how changes can be made, such descriptions can provide information about where systems should not be changed (Halverson, 1994), either because the co-ordinating properties of the representational media cannot be reproduced easily in computerised artefacts, or because change to the processes of work could result in the loss of valuable, task related information.

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 allows us to examine the co-ordination of work *between* agents.

As a final note of interest, one point that is not considered in the analytic framework is how representations come to actually represent information - DC does not attempt to enter into this area, simply considering the *use* of the representations. However the question is not trivial, as the cognitive system must somehow recognise artefacts as symbol bearing representations. The social construction of a representation is necessarily a social process that takes place as various actors come to intersubjectively recognise the meanings within the structure of the representational artefacts. This is, however, not critical to the description of IP activity, and is effectively ignored for the sake of simplicity here. However, it is an area that we believe merits further investigation.

Acknowledgements

This work was partially supported by the CICC project (ACTS project no. AC 017).

References

Brown, J.S. & Duguid, P. (1994) Borderline issues: social and material aspects of design. *Human-Computer Interaction, 9,* 3-36.

Gardner, H. (1985) The mind's new science: a history of the cognitive revolution. Revised edition. New York: Basic Books.

Grudin, J. (1993) Groupware and cooperative work: problems and prospects. In Baecker (Editor) *Readings in groupware and computer-supported cooperative work: assisting human-human collaboration.* California: Morgan Kaufmann. p. 97-105.

Halverson, C.A. (1994) Distributed cognition as a theoretical framework for HCI: don't throw the baby out with the bathwater - the importance of the cursor in air traffic control. *COGSCI Tech Report 94-03,* University of California, San Diego.

Heath, C. & Luff, P. (1991) Collaborative activity and technological design: task coordination in London Underground control rooms. In *Proceedings of the 2nd European Conference on Computer Supported Cooperative Work* p. 65-80. Eds. Bannon, Robinson & Schmidt., Amsterdam, The Netherlands. September 25-27.

Hutchins, E. (1995a) *Cognition in the Wild*. Bradford: MIT Press.

Hutchins, E. (1995b) How a cockpit remembers its speeds. *Cognitive Science, 19,* 265-288.

James, W. (1890/1950) *The principles of psychology. Vol.1.* Edition: Authorized ed. London: Constable.

Larson, J.R., Christensen, C (1993) Groups as problem-solving units - toward a new meaning of social cognition. *British Journal Of Social Psychology, 32,* 5-30.

Newell, A. & Simon, H.A (1981) Computer science as empirical enquiry: symbols and search. In *Mind Design: philosophy, psychology, artificial intelligence* Ed. Haugeland. London: MIT Press. p. 35-66.

Norman, D.A. (1986) Reflections on cognition and parallel distributed processing. In McClelland, Rumelhart, & the PDP Research Group (Eds.) *Parallel distributed processing: explorations in the microstructure of cognition. Vol.2, Psychological and biological models.* p. 531-546. Cambridge, Mass.: MIT Press.

Perry, M.J. (1998) Distributed cognition and computer supported collaborative design: the organisation of work in construction engineering. Unpublished PhD Thesis. Department of Information Systems and Computing

Pylyshyn, Z. (1984) *Computation and cognition: towards a foundation for cognitive science.* Cambridge: MIT Press.

Rogers, Y. & Ellis, J. (1994) Distributed cognition: an alternative framework for analysing and explaining collaborative working. *Journal of Information Technology, 9*, 119-128.

Robinson, M. (1993a) Design for unanticipated use..... *Proceedings of the third European Conference on Computer-Supported Cooperative Work.* Eds, De Michaelis, Simone & Schmidt. September 13-17, Milan, Italy, p. 187-202. Kluwer: Netherlands.

Rumelhart, D.E., McClelland, J.L. and The PDP Research Group (1986) Parallel distributed processing: explorations in the microstructure of cognition. Vol.1, Foundations. Cambridge, Mass.: MIT Press.

Salomon, G. (1993) No distribution without individual's cognition: a dynamic interactional view. In Salomon (Ed.) *Distributed Cognitions.* USA: CUP. p. 111-138.

Simon H.A. (1981) *The Sciences of the Artificial.* 2nd Edition. USA: MIT Press.