

This appeared in Proceedings of the European Cognitive Science Society (1999). Formal citation: Distributed Cognition, Coordination and Environment Design.

Distributed Cognition, Coordination and Environment Design

David Kirsh
Dept. of Cognitive Science
Univ. California, San Diego
La Jolla, CA 92093-0515
+1 858 534-3819
kirsh@ucsd.edu

Abstract

The type of principles which cognitive engineers need to design better work environments are principles which explain interactivity and distributed cognition: how human agents interact with themselves and others, their work spaces, and the resources and constraints that populate those spaces. A first step in developing these principles is to clarify the fundamental concepts of environment, coordination, and behavioural function. Using simple examples, I review changes the distributed perspective forces on these basic notions.

Keywords

coordination, distributed cognition, environment, function of action, interactivity, workflow

Introduction

What makes an environment a good place in which to work? Cognitive engineering would be considerably easier if we had a principled answer to this question. Designers could then consult the book of cognitive and interactive principles and, like civil engineers who build bridges, they could creatively apply principles to build environments that make life easier. At present, no such book exists. And the prospects of ever writing one depends on our understanding a collection of issues broadly related to interaction and distributed cognition which we are just beginning to appreciate. In this paper, I offer a small gesture of appreciation. I will focus on three related issues.

1. **Environment:** what constitutes a work environment? Newell and Simon taught us that an environment cannot be defined independently from the tasks an agent sets out to do. But in their models, tasks typically do not require sensitivity to background social and cultural constraints. Real work environments are a complex superposition of social, cultural, cognitive, and physical constraints. Moreover, the tasks Newell and Simon chose to

- study do not involve dense reciprocal interactions typical of the way we deal with complex machinery such as airplanes, cars, etc. How are we to make sense of the environment as a complex analytical construct when our goal is environment design?
2. **Coordination:** the view I take is that whenever we engage in a real world activity or task, our effective environment is a shifting coalition of resources and constraints, some physical, some social, some cultural, some computational (involving both internal and external computational resources). When this shifting coalition of resources and constraints is appropriately coordinated, the tasks we set out to achieve are accomplished. Lack of coordination leads to lack of success. What does coordination mean? What is involved in coordinating resources and constraints? What are examples at both the social and individual levels?
 3. **The Function of Action:** Do we really know the purpose(s) of our activity? The dominant assumption in psychology, even today, is that the point of activity is to change the environment in a manner that leads, or ought to lead, to goal satisfaction. This view restated in the language of problem solving is that actions should, wherever possible, advance an agent to states in the task environment that lie closer to the goal state. Communicative acts seem the primary acknowledged counterexample to this generalization. Yet if we look at the fine grain of activity it is apparent that action can serve many other functions than the pragmatic one of goal advancement or the communicative ones of exchanging thoughts and accomplishing performative acts. In particular, people undertake actions to save attention, memory and computation; they recruit external elements to reduce their own cognitive effort by distributing computational load. This makes sense only if they are closely coupled to the world. If this observation is accepted it has important ramifications for design. Not only must we design environments to make it easier to take the next step toward goal accomplishment; we need to design them to make it easier for us to perform our other actions: easier to perform epistemic, complementary, coordinative actions. How should the truth of cognitive coupling alter the way we think about design?

What is our activity environment?

From the standpoint of distributed cognition, success in a task, particularly a task involving several participants, is a shifting coalition of agent-environment constraints and resources. Take a familiar case of two people cleaning up after dinner. Major jobs are apportioned by agreement or convention. I'll wash, you stack and tidy up. Then the realities of the situation along with the shared understanding of the participants leads to moment to moment decisions about who will do exactly what and how. To get started perhaps I'll help stack, certainly I'll get the dirty ovenware. But if the pots and pans are fairly large we'll need to cooperate about where they are to be stored in the interim, and where the dishes are to be stacked. Some of this can be done without language since the

constraints of the space beside the sink may substantially determine what goes where. But typically some coordination relies on explicit language use and signalling.

Exactly what is the environment in this example? What are its properties? Do we include the cues and constraints we encode in the environment as we proceed? Do we include our knowledge of our previous history together in cleaning up many times before? Do we include our mutual expectations, our understanding of norms, social niceties, our knowledge of the cost of certain kitchenware? All these things are relevant. But in which analytical construct do we place them: the environment, or the agent?

Normally, one thinks of the environment as a brute presence. It is what is there, outside our skin, right now, and it is the place where things happen and goal relevant states change. It is logically independent of me and my history of interactions, though these can improve my performance because of what I have learned. Thus because of my prior knowledge (either with this environment or elsewhere) I may recognize that a certain wine glass is crystal and likely fragile and costly. So I ought to place it in a safer location and handle it with care. But whatever layers of interpretation I project onto the environment, surely the environment must have the structure to support my projections? And whatever knowledge and agency exists in an agent-environment setup it is to be found within the skin of the agent, the place where executive control resides and knowledge is stored.

With the development of ideas about interactivity and distributed cognition, the notion of the environment as a static external structure devoid of agency and knowledge has come under question. Several issues are at play. First is the question of whether the environment begins directly where the skin of the subject ends or whether the subject's mind should be understood as somehow encompassing a variety of props, aids and controlled external processes over which the subject exercises tight causal control.

Consider guitar players. Some prefer to grow their nails to help their plucking, others prefer to use a pick or plectrum. Is a pick, when attached to a finger for the duration of a performance, external to the player whereas his nail is not? Both can be easily removed after the performance. Both can be ignored by the subject during the performance (unless they break). It seems arbitrary that one grows naturally whereas the other is an artefact that must be fixed in place. Yet if a pick is seen as part of the agent (for the duration of the performance) how can it be part of the environment (for the duration of the performance)? Or what about counting out loud? It is easy to show that by uttering out loud, or by pointing to items, people count better, especially if they are interrupted. We have all had the experience that if someone calls out numbers while we are trying to count there is a real danger that we lose our place or miscalculate. A knee-jerk remedy is to count out loud with enough volume to drown out the unhelpful influence. Are

these external counting noises part of the environment, or do they fall within the boundary of the subject? On the one hand they are literally external to the agent. On the other hand they seem so much a part of the agent's control system that it seems arbitrary that they are externalised versions of inner speech rather than utterances within the internal phonological loop.

It is by now fairly obvious that where we draw the boundary between agent and environment depends on the nature of the explanation we wish to provide and its level or focus of analysis. In most careful accounts of action, behavior can both be viewed and explained differently depending on the questions one is trying to answer and the resolution one is concerned with. The answer to why an action occurred is usually different than the answer to how it transformed the environment. The two are related: achieving ends, the answer to why it was performed, is one way of characterizing what an action is, the answer to how it transformed the environment. But discussion of ends (the why's) allows reference to the thought, beliefs about other ends, and the mechanics of reasoning, which are out of place in characterizing the mechanics of behavior (the how's), since presumably behavior occurs in the environment while thought etc. occurs in the agent. However, once we question the boundary of the agent, the distinction between what is involved in the mechanics of behavior and what is involved in the mechanics of thought becomes more arbitrary. If an agent thinks out loud, is that externalisation part of thought or part of behavior? The answer depends on the pragmatics of explanation. What is to be explained and why? The same applies to using a plectrum when playing a guitar. If the focus is on how a particular sound is achieved we may be interested in how a plectrum is used, but substantially in the same way as we would be interested in how a finger nail would be used. From the standpoint of our explanatory concerns nail and pick are functionally identical. Of course, we could shift focus and ask how the two differ in fine control. But that is a refocusing of the question.

This concern for the pragmatics of questioning is key to understanding the structure of distributed cognition explanations. C.S Peirce used to say that a chemist reasons as much with his hands when manipulating test tubes and glass as with his brain. Presumably he meant that certain manipulations produce meaningful states which can serve as moves in a reasoning sequence. The chemist thinks with his hands. He need not replicate the meaning of that state explicitly in his head if he can rely on it to set the stage for the next meaningful event. Reasoning can be distributed over states inside and outside the head.

Another historical figure who blurred the distinction between what is inside the mind and what is outside was J.J. Gibson (1966). He used to talk about the senses as a system which included the muscles of the eyes, neck, trunk and legs, as well as the lens and other purely mechanical-optical elements. This could even be extended to the microscope used by a microbiologist. For in using a microscope to view a specimen she relies on the microscope as part of her sensory system since perception requires coordination between focusing the

microscope, moving the eye and head, and manually manipulating the slides. To get the right dynamic sampling of the microscopic world requires closing the control loop involving hands, eyes, head, slide and microscope. Where you set the environment and where you set the locus of agency - of control -- depends on the focus of explanation.

This concern with explanatory focus can lead to rather surprising variations in the notion of an agent's spatial environment of action. Imagine, for instance, an oceanographer sitting on the deck of a ship controlling a robot submarine a thousand feet below the water's surface. The robot submarine has pseudo-hands, eyes, and other sensors whose orientation and position can be controlled from on deck. Since the oceanographer has both perceptual and behavioral capacities below, the submarine is acting much like a remote prosthesis. Where is the environment of action?

You might think that the answer is easy: as a question about human-machine interaction the focus is on the point of contact up on the ship between the oceanographer and his instruments; as a question about strategies of marine exploration the focus is on what he is seeing and doing down there. But this is a simplification. In designing computer interfaces, a key focus is on transparency - letting the semantics of action speak through the interface, so that the oceanographer feels he is down there, in immediate causal contact with the ocean floor. Achieving this transparency naturally requires focusing on how ocean floor information is represented proximally, but it also relies on the oceanographer understanding the effects of his behavior on the ocean floor itself, since the perceptual strategies he uses likely presuppose a close causal coupling between himself and the 'optic array' of the ocean floor itself. What counts are the effects which actions have in the domain that matters. So it is essential that interface designers understand the sorts of actions that emerge as the oceanographer interacts with the ocean floor.

Such musings on the nature of our environment of activity have become more pressing in our modern electronic world. The modern white color workplace is a complex knowledge environment in which the flow of information is mediated by an ill understood assortment of technologies, representations, at-hand resources, and shifting teams of people. People engage in many tasks simultaneously, often in ways that cause interference; they interact with each other and with their tools in little known ways; and their primary work space is not confined to the physical region within arm's reach, but is a distributed cluster of 2D and 3D spaces near key resources, computers, telephones and bookcases. Modern workspaces certainly include virtual spaces -- customized computer 'desktops' and applications that have their own worlds of organizational structure, information space, and workflow requirements. Perhaps most significantly of all, the real world is a place we inhabit rather than visit. We live here and return to it. This means that the environment we react to is always a function of the environment we partially created by our own previous actions. Historical properties of

environments are important. As adaptive creatures, we can be sure that we have developed powerful ways of relying on these historical properties to make our activity easier.

This way of looking at the environment complicates the classical analysis of the task environment as the construct in which agents perform their tasks. When Simon formalized and adapted Weber's notion of an institutional role to make it work for task oriented problem solving he emphasised that only a tiny fraction of the properties and events occurring in the environment at large were relevant to the problem solver and so part of the task environment. He laid great emphasis on the distinction between the environment as the activity space in which actions take place and consequences accrue and the problem space, which is the task environment as conceptualised or represented by the subject. The task environment is a construct that applies to the outside world, the problem space is a construct which is supposed to have psychological reality and refers to mental representations inside the agent's head. Two agents could be assigned the same task, and so inhabit the same task environment, but in principle they could create different problem space representations, which might then explain their differential performance. Two agents assigned different tasks, meanwhile were assumed to operate in different environments, even though in other respects those agents were similar. Thus if twin brothers were placed in front of the Tower of Hanoi puzzle and each was told a different set of rules, the two would be assumed to be facing a different environment - a different task environment - since the goals, moves, states, and consequences of actions would be different. In figure 1 this idea of an objectified task environment is represented.

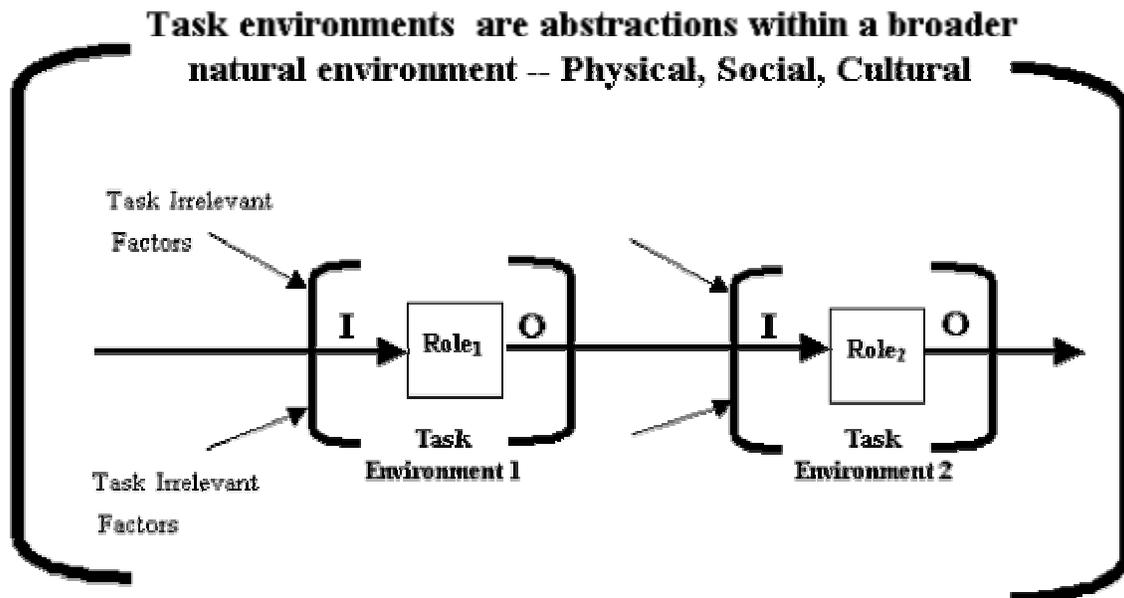


Figure 1. A task environment is a theoretical construct which is supposed to supply the goal relevant consequences of allowable moves in performing the task. Moves or actions are considered allowable if they can possibly advance an agent closer to his or her goal. The point of this illustration is to emphasize that people really operate in much broader activity

spaces than Task environments. We can pretend that they can be viewed as role players, where their only concerns are those to do with performance of their task narrowly conceived. But this is an idealization that eliminates from consideration many of the surprising ways people have of using the environment to help them control activity.

The idea of a task environment was a great analytical advance over previous way of thinking about problem solving, for it meant that there was an external invariant against which performance could be measured. It also meant that the space of potentially relevant actions in the environment was circumscribed. For instance, from a strict task environment perspective there would be no value in breathing since this was not a task relevant action. Breathing is a background condition of some relevance, but it lies outside the boundary conditions of the task, since it cannot in principle advance the agent toward task completion. The same applies to talking out loud, using a pocket calculator to simplify computations and so on.

Not surprisingly, there is a way of seeing these phenomena as rational within the task environment/problem space paradigm. Using a calculator, or repeating ideas aloud, for example, can themselves be seen as ways of dealing with a certain problem. In the case of calculating, the task might be to solve a sub-problem that is otherwise quite difficult to do in the head or on paper; in the case of talking aloud, the task might be to improve working memory. So the formalism is flexible enough to cope with many complaints. Indeed, the task environment notion is rich enough to accommodate cultural and social factors affecting cognition, since these may be interpretable as constraints on permissible goals, sub goals and actions.

But despite this flexibility the basic orientation of pitting the individual against a task, and seeing this relation as solving a problem, is one which takes the key elements of a solution to be those that occur inside the agent's head where the steps in a problem space are actually made. The spirit of the problem space approach is not to ask how problem solving is distributed over environment and agent, although it is possible to ask this question. See, for instance, Simon and (Larkin and Simon, 1986), and (Larkin, 1989). The spirit is to be individualistic. Thus although the formalism allows us to ask why one given external representation is better than another, it remains an essentially one person approach to the problem. It does not provide the wherewithal to explain how new and better representations are created through interaction with other people, particularly if, somehow, there is an emergent ingredient not present in the problem space of any group member. Moreover, the formalism offers little or no help in letting us understand why people develop such surprising ways of using the environment to help them control activity.

In short, insofar as distributed cognition takes a more group orientation to problem solving and activity management, it operates from a different mind set than what is now the classical approach to problem solving and it provides a different set of explanatory tools.

Coordination

A basic tenet of distributed cognition is that although humans are key players in a coordinated system of distributed influences, they are not the only influences. When I drive to work, for example, I end up at my destination only if my car behaves as it should and the other cars behave as they should. In an odd way, my car and I - and to a lesser degree the street and traffic -- are partners in getting me to my destination. My car and I are co-dependent. If this seems an unnatural way of speaking at present that is because of the apparent command and control nature of humans' relationship with cars. Cars are there to serve our needs. They are not partners with decision making rights. But as cars become increasingly computerized it will seem less odd to think of our sharing control in the job of getting us safely to our target. It is not that as drivers we humans may expect to lose primary agency over our cars. It is that at finer levels of detail, key 'decisions' about how our car behaves are, and will increasingly, be made by the car itself.

Coordination is a technical term for this sort of partnership between humans and the resources and constraints in their environments. What do we know about coordination? What are examples of coordinative mechanisms and processes?

The first thing to say about the concept of coordination is that it applies at several different temporal resolutions. At a coarse grain we can ask about the long term physical setup, about the enduring lines of communication, the institutional structures in place to help us work. All these structural elements of the set-up are important determinants of how easy it is to achieve effective coordination between ourselves and our environment. At a finer level of detail, we can ask about the short term coalitions between people, about the temporary partnerships they make in helping each other to complete a task. At a finer grain still, we can ask about momentary modifications to the environment that help us to communicate, to reason, to perceive, to interpret.

A second important aspect of coordination is whether it is achieved through explicit - usually symbolic - means, or whether it is achieved through implicit - usually non-symbolic means. The distinction is simple enough. When four people attend a work meeting in which work plans are discussed and represented on a whiteboard, or encoded in a project plan, they are engaged in an explicit coordinating activity - the meeting - and they rely on a symbolic device - writing on the whiteboard - to help coordinate the meeting itself, and to help coordinate the project. The same applies to chess playing. Coordination requires taking turns. If someone violates this explicit convention they are unambiguously told. By contrast, when animals develop trails through the forest, or people develop paths through snow, they are not explicitly coordinating their activity. Trails emerge, the way giant termite hills emerge. Locally optimising behavior leads to global configurations without explicit communication between participants, and without symbolic communication. There are no explicit mechanisms of

coordination. No social conventions, no discussions, no maps, charts, way finding devices or way finding representations.

Let us start with some simple cases of explicit symbolic coordination among people. What are some enduring ways teams of people explicitly coordinate their actions to achieve complex outcomes?

One of the simplest ways of explicitly coordinating activity is through production lines. The idea behind a production line is that participants are assigned very specific roles using very specific machines. Team members have few if any degrees of freedom. Outputs of one role are the inputs to another. The whole process, and the expectations about team members performance are explicitly documented. See figure 2.

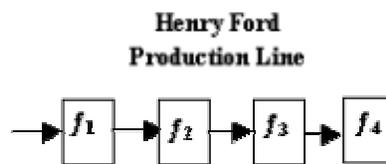


Figure 2.

A less simple way of achieving coordination but no less explicit can be found in orchestras. Orchestras partly resemble production lines in that team members have specific roles to play. However they differ in several interesting ways. First, activity is simultaneous, so the output from one player is not input to another; rather all outputs are part of a common emergent whole. Second, members' role varies with piece, so members must be 'rapidly reprogrammable'. Third, members are encouraged to listen to the entire groups' performance rather than be myopic and solely concerned with their own performance. Fourth, reading a score requires significant interpretation on their part, and so leaves room for considerable individual variation. And fifth, there is, of course, an orchestral leader, a conductor whose job is to explicitly coordinate performance.

The score is an explicit symbolic representation of role. The physical arrangement of chairs, players and conductor is an explicit non-symbolic coordination. If each player's interpretation of his or her score were sufficiently constrained there would be no need for a conductor. Roles could be counted on to 'sum' correctly, to produce an emergent whole in accord with the composer's intent. But in fact, orchestral scores underconstrain performance. They say nothing of how players should adapt to the acoustics of a hall, or to changes in orchestral size. They are mute on whether the violins should sound louder than the cellos when both are playing forte. Even worse, they specify tempo in qualitative terms such as *andante*, *lento*, *allegro*, and they use subjective modifiers such as *allegro con molto* to convey mood. A central conductor, therefore, is needed to help coordinate tempo, coordinate feeling, balance volume, and in general, to blend the emergent sound in a way that is, in his or

her expert judgment, musically coherent and beautiful. A conductor serves the purpose of communicating global constraint to all locally constrained participants.

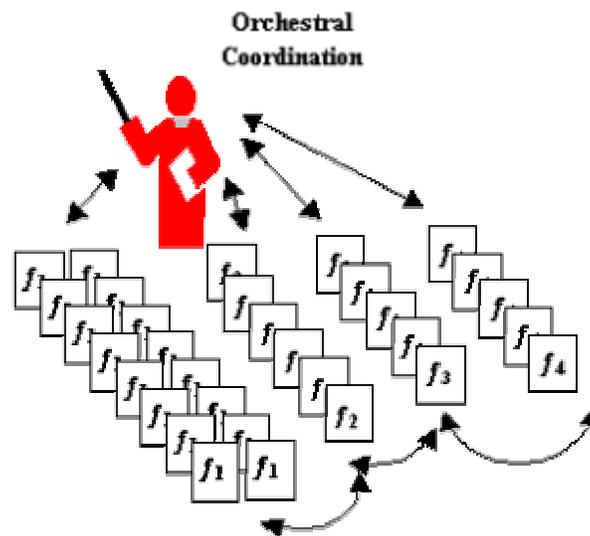


Figure 2. An orchestra leader coordinates the activity of a team of players who each have a pre-assigned role, characterized here by the function symbol f_i .

Another example of explicit coordination can be found in American football, where the quarterback plays a different leadership role in coordinating team activity. In American football, players share a huddle where the quarterback calls a play and offers advice, reminders and commentary. Once the play is announced, however, all players know their role for the next phase of activity. They have practiced it many times before and know what they are to do relative to one another. If they execute these roles correctly a pattern is created that wins them yardage.

As with musicians, though, their role is not specified in complete detail. Because football is a competitive sport, situations dynamically change on the field as a result of opponents' activity. All players, therefore, also need to understand the point of their role vis. a vis. the play as a whole. This helps them to know how to modify online their activity to make it contribute to the spirit of the play. If things go their way, this network of local decisions by players leads to globally desirable outcomes. Football and orchestras are examples where coordination emerges because a central figure adds global perspective to a system of decentralized but rule governed players. A leader is required to add constraint, but the building blocks for success lie in the team members' roles. In the music world these roles are represented in sheet music, each player receiving his own copy. In football individual roles are typically represented as part of the group's activity and displayed on whiteboards, discussed on the field, and observed on video. A team

leader can be an effective mechanism of group coordination if the coordination problem satisfies at least two conditions:

1. it is possible for the team leader to track (represent) overall group behavior and classify it meaningfully into better or worse states; and
2. it is possible to communicate with all participants as needed without creating a bottleneck. That is, it is possible to easily broadcast group wide advice or possible to target specific advice in time to be useful.

In cases where either of these conditions fails it is necessary to rely on flatter, less hierarchical organizations to achieve coordination.

Scheduling is an interesting case which demonstrates the relevance of these two conditions. The problem of scheduling all the meeting times for a group of participants is an example of a computationally intensive problem. In its general form it requires finding an allocation of meetings to times so that everyone can attend all their meetings without conflict. Meetings may be with the whole group or a fraction of the group. Since the problem in its worst case requires trying out all possible times for all possible groups it cannot be solved any faster than by exhaustively enumerating all combinations of times, and meetings. Put in a more interesting form, to say that the problem is computationally hard is to say that there is no good way of representing who is meeting when such that the overall goodness of a set of meeting arrangements can be easily read off. Any descriptively adequate representation will be deceptive in at least one case, leading the scheduler to believe that he or she is close to a solution when in fact they are not.

Scheduling cannot be done well by a single point person. The reason, again, is that the team leader cannot reliably track the goodness of meeting arrangements. A shift of one meeting can result in an even worse allocation, and there is no easy way of telling that in advance of trying it. Of course, a quick solution could be had by broadcasting all possible arrangements and listening for the ones that generate no conflicts. But this would create an impossible demand on group members as well as a communication bottleneck. Since both conditions one and two fail for the problem of coordinating schedules, coordination using a more decentralized method ought to be better than using a centralized coordinator (if there is a better method of coordination at all).

Can it be done any better in a decentralized manner? In principle, it cannot, since it is an NP problem and so cannot be solved noticeably faster by increasing the number of people trying to solve it. Nonetheless, in practice, it is desirable to spread the computing load around by finding a representation that all participants can interact with, and coordinate the time at which they interact with it, because people will find work arounds. Rather than blame a group coordinator, group members who are jointly involved in solving the problem tend to relax the problem by moving some of their meetings out of the problem range. Strictly

speaking this violates the conditions of the formal problem, but it does solve the coordination problem because it allows participants to decide for themselves which meetings are most important - an item of information that would have just increased the complexity of the group coordinator's job. Often this extra information is added by annotation to a central representation, leading to an emergent solution that would be hard to achieve centrally.

Coordination, as I have been describing it so far, has relied on some form of explicit symbolic communication. An important tenet of interactive and distributed cognition is that coordination can be achieved in implicit and non symbolic ways too. Here are two such cases that also show how it is possible to achieve coordination between team members by methods other than leadership, sharing central representations, and role playing.

The first case of non symbolic coordination shows up at a fairly low level of physical coordination between team members. Imagine three groups of untutored singers setting out to sing a new song in three different rounds. Having just returned from such an event I can report anecdotally that all groups fell out of phase quite soon and collapsed into an in phase form where everyone was singing in unison. The explicit role of each team was to begin singing at a particular point in the other's song ,and thereafter maintain a constant difference. The implicit coordinating force, however, was an attractor to sing in unison.

In dynamical systems containing oscillators, it is common to find the overall system falling into certain global rhythms. This preference for one rhythm or frequency over others is called entrainment. It happens when the timing of repetitive motions by one oscillator influences the motions of the other so that they couple, and their phases lock together. The tendency of groups of walkers to end up marching in step is another example where their actions become coordinated without intentionally seeking that coordinated state. Obviously this coordinating force could be explicitly harnessed to produce predictable effects. In such a case, we would have explicit but non symbolic coordination.

Other cases of phase locking in biology have been documented by Glass and Mackey (1988) in their discussion of animal gestures. Among humans a further example is found in students learning to play bongo drums. To play an extended roll on the bongos it is necessary to coordinate hand motion so that when one hand goes down the other goes up. If you have tried this yourself you will know it is not easy to maintain a roll at high speed. One hand tends to capture or entrain the other and both hands slip into a zero phase lag where they hit the drum head simultaneously. Practice can overrule this tendency, but there is no denying that it is a natural state. Entrainment, then, is a form of coordination that does not involve symbolic communication between participants.

The second and last case of non symbolic coordination I will cite is drawn from my own recent experience. For the last few months my mother in law has been

staying with us. It is evident that she has different notions of what it means for the house to be in order; in particular, she has a different idea of the state one should leave the kitchen in. Because of her influence on the household our behavior has settled on a new equilibrium state - one in which all dishes are put away, either in cupboard or dish washer. Now why did this happen? She has not requested this explicitly, although it is obvious that she approves of it. So undoubtedly expectations have played a role in our temporary conversion. But on a more theoretical level I think it has happened because some states of kitchens are more powerful attractor states for groups than others. If you assume that you will continue to work and cook in that kitchen it is often better to put dishes, cutlery and ingredients away. Hammon et al.(1995) discussed how it is desirable to return dishes to the kitchen, placing them in the sink, dishwasher or cupboard rather than leaving them in the living room, or dining room where they were last used, because if the dishes are in a known place, especially a place near to where they will be used subsequently, it simplifies planning by cutting out the stage where you must first hunt for crockery. The next step, which involves placing all dishes in the dishwasher shortly after their use, is also justified if the increase in clutter which another person creates causes more inconvenience than placing the dishes in the dishwasher right away. Coordination, or rather a new state of coordination, was achieved in our kitchen because our overall performance was enhanced by putting the dishes away. It represents a gain in efficiency.

Now my reason for discussing a few forms of coordination is that when viewing agents acting in their environments of activity we believe it is essential to understand how they coordinate their own activity with the other causally important elements of the environment. When other agents are present, the motive for talking about distributed cognition is usually that cognition is thought to be distributed across several minds. In fact, though, distribution refers to the close causal coupling between all the causal influences driving the agent-environment system to its goals states. The mechanics of coordination become the focal point of explanation.

As designers our objective is to understand the mechanics of coordination well enough to add structure to the environment - scaffolding, cues, prompts, artefacts and communication devices - to make work places more effective. In extreme cases we might completely redesign the spatial layout of the workplace and the basic workflow. My discussion so far has been rather abstract. Instead of examining actual mechanisms of coordination -- mechanisms such as day planners, agendas for meetings, market systems, collaborative software, annotations, whiteboards and the like - I have considered how citing these mechanisms can help to explain coordination. So for instance, white boards, day planners, agendas and so on, serve as mechanisms for adding constraint to a distributed system. They are not just resources to be used the way a calculator is a resource to be used -- that is, as a method for offloading memory and computation -- they figure as coordinative structures.

The Function of Action

So far we have been considering coordination as it emerges among groups of people and their environment of action. In the majority of my own empirical work I have focused primarily on the coordination that emerges between a single individual and his or her environment of action. (Kirsh and Maglio, 94), (Kirsh, 95,95a) The two - person and environment -- are coordinated in the sense that reaching a goal state depends on both sides doing their part. What I have been continually surprised by in these studies is the diversity of ways we have of recruiting the environment as our cognitive ally. This too is a form of coordination, for we coordinate our internal processes with external ones in a tightly coupled manner. In observing everyday behavior it was possible to find a variety of actions that were not of immediate pragmatic value to agents. They were reliably present, and make good sense given an organism's enduring presence in an environment, but they are not performed to bring the agent closer to some pragmatic goal defined as a certain state of the environment I will close this paper with a simple list of a few examples of these non-pragmatic actions that show up in the phase of activity occurring before people have decided what to do. Such non-pragmatic actions appear in all phases of activity. Removing clutter, for instance is one such action that may occur in any phase of activity. Here, though, I will mention only those that appear early on. They cover a range of temporal duration: from 100 millisecond range as in Tetris, seconds in scrabble and jigsaw, minutes in cryptarithmic, geometric problem solving, and book organizing.

During scrabble, people reshuffle their tiles as if to facilitate, or cue, recall of possible words.

- While searching for an appropriate jigsaw piece in solving a jig saw puzzle, players tend to construct groupings such as corner pieces, edge pieces, same color, similar shape. These intermediate steps aid visual search, but their function is cognitive or epistemic, in that they do not actually bring players physically closer to their pragmatic goals.
- When setting out to organize a bookshelf according to subject heading, the sort routine often followed involves distributing books first on the floor or in different regions over the bookshelf, as if to prove that a particular subject arrangement makes sense. These interim steps frequently require revision: their function is as much cognitive or epistemic as practical, since the early arrangements of books may not find their way into the final arrangement.
- When playing Tetris, players have little time to choose their target placement, yet they rotate pieces often four or five times more than necessary. This extra rotation is not evidence of flailing, but rather plays a functional role in their computation of their goal placement.
- In solving cryptarithmic problems subjects often mutter, write down intermediate conjectures, partial results, rewrite elements in different

places on their scrap page. These actions seem to help problem solving, although it is not always evident how.

- In solving simple geometry problems subjects try out a range of different constructions. The function of constructions is not necessarily to permit solution directly, but often to make certain properties of the structure more explicit, or to prompt the agent to notice similarities between this problem and others seen before.

In each of these cases, agents recruit aspects of their environment, or create structures in their environment that link with internal states in creative ways. When the timing of these structural additions is appropriately related to internal processes improved performance results. This interplay between creating structure in the environment and projecting structure onto the environment (interpreting results) is a key process in agent environment coordination.

Although the examples just mentioned contain no words of analysis, research into interactive and distributed cognition is highly detailed. It takes into account broader descriptions of subjects' experience and their history of interaction with similar items, as well as measures of timing. As our workplaces continue to integrate digital enhancements it seems obvious that we have the opportunity to increase the sophistication of the coupling between agent and environment. It is my belief that studies of interaction and coordination will become increasingly central in this effort.

Acknowledgements

I thank the National Institute of Aging, grant number AG11851 and the National Science Foundation grant number NSF nnnn, for support of this research.

References

Gibson, J. J. (1966). *The Senses Considered as Perceptual System*. Boston: Houghton Mifflin.

Glass and Mackey (1988). *From Clocks to Chaos: The Rhythms of Life*. Princeton University Press, Princeton, NJ.

Hammond, K.J., T.M. Converse and J.W. Grass. (1995) The stabilization of environments. *Artificial Intelligence* 72: p. 305-327.

Kirsh, D. (1995a). Complementary Strategies: Why we use our hands when we think. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Kirsh, D. (1995b). The Intelligent Use of Space. *Artificial Intelligence*. 73: 31-68

Kirsh and Maglio, (1994), On distinguishing Epistemic from Pragmatic Actions, *Cognitive Science*. 18: 513-549.

Larkin, J. H. and Simon, Herbert (1987). Why a diagram is (sometimes) worth 10,000 words. *Cognitive science*, 11. 65-100.

Larkin, J.H. (1989). Display Based Problem Solving. In Klahr and Kotovsky (eds.), *Complex Information Processing*. Lawrence Erlbaum. 1989.

Port Robert, Keiichi Tajima, and Fred Cummins (1998). Speech and Rhythmic Behavior in G. J. P. Savelsbergh, H. van der Maas and P. C. L. van Geert (eds), *The Non-linear Analyses of Developmental Processes*, Royal Dutch Academy of Arts and Sciences, Amsterdam, 1998.

Other Articles

Kirsh, D. (2000). [A Few Thoughts on Cognitive Overload](#). *Intellectica*,

Kirsh, D. (1999). [Distributed Cognition, Coordination and Environment Design](#), *Proceedings of the European conference on Cognitive Science*

Maglio, P. P., Matlock, T., Raphaely, D., Chernicky, B., & Kirsh D. (1999).

[Interactive skill in Scrabble](#). In *Proceedings of Twenty-first Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum.

Knoche, H., De Meer, H., Kirsh, D. (1999). [Utility Curves: Mean opinion scores considered biased](#). *Proceedings of the Seventh International Workshop on Quality of Service*

Kirsh, D. (1998). [Adaptive Rooms, Virtual Collaboration, and Cognitive Workflow](#). In Streitz, N., et al. (Eds.), *Cooperative Buildings - Integrating Information, Organization, and Architecture*. Lecture Notes in Computer Science. Springer: Heidelberg.

Elvins, T, Nadeau, D., Schul, R., Kirsh, D. (1998). [Worldlets: 3D Thumbnails for 3D Browsing](#). *Proceedings of the Computer Human Interaction Society*.

Kirsh, D. (1997). [Interactivity and MultiMedia Interfaces](#). *Instructional Sciences*.

Elvins, T, Nadeau, D., Schul, R., Kirsh, D. [Worldlets: 3D Thumbnails for Wayfinding in Virtual Environments](#) *UIST97* 1997.

Kirsh, D. (1996). [Adapting the Environment Instead of Oneself](#). *Adaptive Behavior*, Vol 4, No. 3/4, 415-452.

Kirsh D. (1995). [Complementary Strategies: Why we use our hands when we think](#). In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Kirsh, D. (1995). [The Intelligent Use of Space](#). *Artificial Intelligence*. 73: 31-68

Kirsh, D., & Maglio, P. (1994). [On distinguishing epistemic from pragmatic action](#). *Cognitive Science*. 18, 513-549. [ps file, 1000K]

Kirsh, D., & Maglio, P. (1992, March). [Perceptive actions in Tetris](#). In R. Simmons *AAAI Spring Symposium on Selective Perception*. [ps file, 191K]

Kirsh, D., & Maglio, P. (1992). [Some epistemic benefits of action: Tetris, a case study](#). In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum. [ps file, 875K]

- Kirsh, D., & Maglio, P. (1992). [Reaction and reflection in Tetris](#). In J. Hendler (Ed.), *Artificial intelligence planning systems: Proceedings of the First Annual International Conference (AIPS92)*. San Mateo, CA: Morgan Kaufman.
- Kirsh, D. et al. (1992). Architectures of Intelligent Systems, in *Exploring Brain Functions: Models in Neuroscience*. John Wiley.
- Kirsh, D. (1992). PDP Learnability and Innate Knowledge of Language. In S. Davis (Ed.), *Connectionism: Theory and practice (Volume III of The Vancouver Studies in Cognitive Science, 297-322)*. NY: Oxford University Press.
- Kirsh, D. (1991). [Foundations of artificial intelligence: The big issues](#). *Artificial Intelligence*, 47, 3-30.
- Kirsh, D. (1991). [Today the earwig, tomorrow man](#). *Artificial Intelligence*, 47, 161-184. Reprinted in M. Boden (ed) *Philosophy of Artificial Life*. Oxford University Press (in press)
- Kirsh, D. (1990). When is information explicitly represented? In P. Hanson (Ed.), *Information, language, and cognition. (Volume I of The Vancouver Studies in Cognitive Science, 340-365)* Vancouver, BC: University of British Columbia Press.
- Kirsh, D. (1987). Putting a price on cognition. *The Southern Journal of Philosophy*, 26 (suppl.), 119-135. Reprinted in T. Horgan & J. Tienson (Eds.), 1991, *Connectionism and the philosophy of mind*. Dordrecht, ND: Kluwer.
-