Adaptive Rooms, Virtual Collaboration and Cognitive Workflow

David Kirsh

Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093-0515, USA
Email: kirsh@ucsd.edu

Abstract. This paper introduces the concept of Adaptive Rooms, which are virtual environments able to dynamically adapt to users' needs, including 'physical' and cognitive workflow requirements, number of users, differing cognitive abilities and skills. Adaptive rooms are collections of virtual objects, many of them self-transforming objects, housed in an architecturally active room with information spaces and tools. An ontology of objects used in adaptive rooms is presented. Virtual entities are classified as passive, reactive, active, and information entities, and their sub-categories. Only active objects can be self-transforming. Adaptive Rooms are meant to combine the insights of ubiquitous computing -- that computerization should be everywhere, transparently incorporated -- with the insights of augmented reality -- that everyday objects can be digitally enhanced to carry more information about their use. To display the special potential of adaptive rooms, concrete examples are given to show how the demands of cognitive workflow can be reduced.

Keywords. cooperative buildings, collaboration, cognitive workflow, cognitive ethnography, ontology, virtual collaboration, virtual environments, virtual objects

1   Introduction

The goal in designing virtual collaborative environments is to allow individuals to do everything they can do in real shared spaces and more. As in real spaces people must be able to talk to one another, move around, make diagrams, build models, highlight points of interest for others to consider, and jointly edit documents or 3-D models. The ultimate promise of virtual reality, though, is that users will be able to do things they cannot do in real life: they will be able to conduct new kinds of scientific, business and social explorations via meetings held in "outer space," within a "molecule," inside the "combustion chamber" of an automobile engine, suspended in the "atmosphere" above planet Earth, or in Ms. Frizzle's Magic School Bus. In such cases, users will be able to jointly interact with simulations. One particular aspect of this interaction we are exploring and will report here is how to design virtual environ-

ments to dynamically adapt to the workflow needs of participants – both `physical' workflow and cognitive workflow. How should we embed simulations, information spaces, and other computational tools into virtual environments to facilitate collaborative activities?

Workflow adaptation is a thorny problem. At the most familiar level, collaborative workflow is understood in a pragmatic or `physical' manner as the activities and sub-activities – the tasks and sub-tasks -- which collaborating partners perform. Any typical job, such as assembling an electric motor, can be decomposed into a lattice of component activities. Parts must be collected, compared and sorted, then aligned correctly and fastened. Because some of these activities must be performed before others there is a partial temporal ordering on the task decomposition, hence the use of a decomposition lattice. In ordinary physical environments, collaboration makes this lattice more complicated because we must also decide who will do what; but the temporal structure of the job remains essentially the same. In this context, it is clear what an adaptive room should do: It should adapt the space, furniture and resources available, to the special needs of each sub task. If the task of comparing the parts requires lighting that is brighter than normal, then when that subtask is being performed the lighting should automatically be brighter. If the task of sorting parts requires special bins in which to group parts, then for the sorting phase new bins should appear as needed. Similarly, if alignment is facilitated by a jig, then a jig should be present to hold or re-orient the main motor. The list of useful adaptations can be extended. If several people wish to help in the assembly, and they have not decided to work separately in assembly line fashion, then the physical space around the main assembly platform should expand to comfortably accommodate more people.

This last adaptation – morphing of walls and furniture -- is one we expect to arise in most collaborative tasks in adaptive rooms. Unlike ordinary physical environments where limited space and chairs around a table or computer screen invariably means that some people must stand, in virtual rooms, any number of avatars can be seated because we can expand or deform the table to accommodate convenient placement of chairs. The computer screen, the whiteboard and bookshelves, the corkboard and the stick'ems, can all adapt. Any facet of the environment that is not currently useful may be temporarily removed. Any facet of the environment that might be useful may be temporarily added. To take another example, if I have been using my office to write an essay on adaptive rooms, and my collaborators on a different topic arrive, it is likely that my messy desk will be an inconvenience to us all. Since I wish to keep the state contained in the arrangement of papers on the table, but I also wish to have the workplace optimally configured for my current collaborative activity, I will either create a new room for this new collaboration, or adapt my office. Because the proliferation of virtual rooms for each collaboration and each activity would soon become disorienting and awkward, a better solution to this problem is to have my books and papers contract to a 3D icon, my bookshelves recede, and the whiteboard expand, all to return automatically when my visitors leave. If social needs require it, extra chairs may be whisked into the room, and any writing pads, markers and related office supplies can be provided as needed.

Adaptation to workflow conceived of as physical task decomposition is a problem designers deal with daily. It may seem, therefore, that although going virtual adds options to the design space, it adds nothing, in principle, to the design problem itself.

That is, whatever workflow requirements there are in physical space, the same requirements still apply in virtual space, though we now have new ways of meeting them. Surprisingly, this is not correct. Virtual environments allow us to create novel task decompositions. For instance, returning to our assembly task, a challenging assembly problem can be tackled by many people simultaneously in ways that are simply impossible in physical environments. Consider the possibility of four people each working with clones of the same parts, all trying to find their own solution.[1] In the physical world, instantaneous cloning is not possible; an object can be in only one place at a time; and any clones that are created soon develop their own histories and separate identities. Changes in one clone cannot be instantly propagated to all the others. The result is that many of the task decompositions we face in the everyday world are the product of physical invariants. Eliminate these physical limitations and the task lattice may be altered.

The same possibility of changing the underlying constraints of a task applies to the cognitive dimension, the cognitive workflow of a task. Cognitive workflow may be conceived of as the changing pattern of cognitive demands placed on an agent as it performs the various component activities of a task. It reflects a task's cognitive decomposition. In the subtasks of assembly for example, there are going to be more or less demanding phases for memory, perception, planning, and so forth. What makes identifying the cognitive workflow of an activity particularly difficult to anticipate, is that agents develop cognitive and interactive strategies that alter the cognitive landscape of a task. Cognitive strategies evolve in partial reaction to the resources that are available. If there is a writing pad nearby, and people have memory tasks, sooner or later they are going to discover the utility of writing down what they have to memorize. Memory can be offloaded. Similarly, people can learn to put their eyeglasses in a standard place, saving them from having to remember where they last laid them down. There are countless techniques of this sort, countless ways people discover of reconfiguring the cognitive costs of a task. Returning to assembly workers, if they have visualization tools which let them simultaneously view an emerging assemblage from different perspectives, the cognitive effort of mental rotation, normally required to determine if a candidate piece would be well placed, can be reduced. Similarly, if workers can peak over their shoulders to see what their colleagues are doing, they need not remember some of their earlier decision rationales. They can simply switch to another's assemblage if it seems more promising, or review how things would have worked out had they themselves pursued a different line of attack earlier. Perhaps they then will `jump' back to an earlier state, quickly undoing their previous decisions. Or perhaps they can be given the equivalent of layers in which to work, thereby making it possible to undo some decisions but leaving others. To be sure, as we create new assembly environments and tools, there will be complex problems of coordination both at an individual and a group level, when several lines of attack prove fruitless and collaborators must decide whose assemblage to pursue and how far to backtrack. These will raise new cognitive demands of their own. But such problems are virtuous, since they offer the possibility of all parties benefiting from parallel search. We cannot predict in advance how much this will reduce the overall complexity of the task, or what the complexity-performance trade-

off will look like. The point stands, however, that in a virtual environment the notion of workflow, both cognitive and `physical' workflow, is not as constrained as in true physical environments. The challenge of designing adaptive rooms, accordingly, promises to be an ongoing one, raising issues never before anticipated.

Given the exciting possibility of altering the structure of a task by altering the `physical' constraints of the environment in which it is performed, adaptive rooms raise the provocative hope that they may take ubiquitous computing – a chief motive for creating adaptive rooms --one step beyond its currently envisioned form. The original idea of ubiquitous computing, (Weiser 1993), is that a single computer should not be the locus of computation in one's home and business environment. Technology should be embedded and distributed in the physical environment in an invisible and transparent manner. In a rich `UbiComp' environment there would be hundreds of computationally driven gadgets or smart appliances throughout, each one part of a larger system of coordinated devices. These objects transmit and receive signals from neighboring objects and often act on them in a context sensitive manner. Many of the objects communicate tacitly, using ambient sensing, such as sonar or video to pick up change. Although these objects do not transmit coded signals to each other they still interact in helpful and often apparently intelligent ways. In particular they are not intrusive. A classic example is a sensor which recognizes my entrance and adjusts the room temperature, lighting, and music, to my preferred levels. As I move about, sensors of this sort will interact with the phone system and help decide which phone to ring and whether to use a distinctive ring for me. If I close my door after a colleague arrives, my room will `know' I wish privacy and route my telephone calls to voice mail. When I sit down at a computer monitor the blinds will draw appropriately to prevent glare. Some of this coordination is achieved through explicit signaling between devices, other aspects of the coordination is achieved tacitly by detecting my movement through ambient sensing.

Our goal in adaptive room research is to unite the power and flexibility of virtual environments with the insight of ubiquitous computing. In effect, it is to simulate the behavior of systems of future smart objects, and to enhance the possibilities of ubicomp rooms by using digitally enhanced objects. To properly design adaptive rooms along these lines we must be sensitive to three requirements.

1. The various cognitive and physical workflows occurring within it;
2. We need to tune rooms to the social needs of users as they interact.
3. We need to maintain environmental coherence across room changes. Adaptive rooms are supposed to be comfortable habitats, not Alice in Wonderland nightmares.

The remainder of this paper has two parts. In the first and largest part, I discuss the ontology of objects necessary for understanding and designing adaptive rooms. In the second, I offer a few illustrations of how cognitive workflow can be changed by the clever use of environmental resources, and a brief example of how these workflow ideas figure in a collaborative task. I then draw a few implications for the design of adaptive rooms.

2   An Ontology for Adaptive Rooms

The highest principle of HCI design holds that it is the environment which should be adapted to users rather than users who should have to adapt to the environment. Characteristically, this has meant organizing the layout and tools available in a software application in a convenient and customizable way to make it easy for users to rearrange their working environments the way they like.   When this principle is extended to more dynamic virtual environments, where environments may automatically adapt to users' activities in more sophisticated ways than simply activating and deactivated tool sets or changing the position of icons, a host of new problems arises that are associated with naturalness of change, plausible adaptation to context change and environmental coherence. If the result of changing a room to accommodate a change in the social context, such as shifting from a discussion to a Powerpoint presentation, is to change the design so radically that the new room bears little resemblance to its previous self then it hardly makes sense to call the newly designed room the same room as the original. It will have been stretched beyond recognition and the cognitive benefits to the user of knowing where he or she is spatially will be destroyed by the confusion that arises from cognitive disorientation.

To properly understand such problems requires empirical study of people's reaction to room adaptation - an empirical study we are just beginning. To be sure, there is empirical precedent for room redesign in the reconfigurable walls which Le Corbusier promoted in the 1930's and 40's, (Le Corbusier, 58), and the open plan education movement in the United States of the 1950's, (Bay, 79), which further developed some of these ideas.   But the physical rearrangement of walls and furniture to meet the educational and activity needs of groups was not an architectural transformation which happened automatically and transparently.   There were no invisible agents and smart objects acting behind the scenes making adaptations in response to the physical task and social cues in the environment. Rather an open plan room was reconfigured only after a sustained process of discussion and negotiation among the participants. Human deliberation was involved. Consequently, the empirical studies of open plan teaching do not carry immediate implications for Adaptive Rooms. Indeed, since the very teachers and students who were using open plan spaces were the ones who physically reconfigured the space, the sorts of concerns about coherence typical of automated adaptation did not arise in open plan classrooms because members had time to adapt themselves to the new layout.

To set the framework for a principled study of Adaptive Rooms it will be useful to begin our inquiry with the ontology of entities that inhabit and constitute them. At a concrete level these include walls, whiteboards, furniture, agents, and other potentially smart or self-adapting entities. But the goal here is an abstract ontology. As a first cut let us distinguish four types of virtual objects: passive, reactive, active and information spaces. Each of these poses a different type of problem for programmers.

2.1   Passive objects

A virtual object is passive if it can change absolute state (shape, color etc.) or relative state (position, orientation) when a human agent or some other active object interacts with it, but is otherwise unaffected by changes in the absolute or relative state of other objects. For instance, we assume that in a simple virtual environment, one which lacks illumination and shading, a simple object, such as a table, will be unaffected by the activity of neighboring objects. Drop a heavy vase on it and the table does not dent. Push a chair against it and it does not move. Humans may rearrange the location of chairs, or possibly alter their color or texture. But such changes do not occur in automatic response to other changes. The primary architectural elements in simple environments -- walls, ceilings, and floors -- are passive in this way. So is fog, and so, of course, have been the early versions of virtual objects.

2.2   Reactive objects

A virtual object is reactive if it can change absolute or relative state, not only as a result of actions on it by agents but in response to changes in other objects. For instance, objects which break when struck by another object are reactive objects because they change their shape and number – their absolute state – and change their position and velocity – relative state – in response to the change in position and momentum of the object colliding with them. We may summarize the difference between reactive and passive objects by saying that reactive objects obey `physical' laws of interaction. Such laws specify how objects cast shadows over surfaces (hence how they become shaded by the other objects); how the absolute state of objects change according to the force, and shape of objects impacting them; how objects change their relative state (position, velocity, acceleration, orientation) as the result of forces. Some of these laws, such as the laws governing reflectance and shading, are global laws of physical interaction and apply throughout a virtual environment. Other laws are highly specific, potentially complex physical constraints, and apply only to particular objects in particular conditions. For instance, the way a couch deforms when a person sits on it depends on the person's weight, speed of sitting, as well as the relative position of other cushions. General physical principles are at play here, so these constraints do not lie outside the physics of the virtual world, but the initial and boundary conditions of pillows on a couch are so unique that it is best to think of the couch as having its own specific laws of interaction -- the physics of couches. As the number of objects increases, it becomes extremely hard to specify a realistic set of interactive physical laws, though luckily, for many interactions, a feeling of realism may be achieved by rough approximation to these laws.

   The items of furniture found in Adaptive Rooms may behave as reactive objects, or they may be more active – closer to what we would call smart furniture – and so capable of self-transforming in ways that go beyond reaction to simple physical events. It is the designer's choice. He or she may choose to program virtual chairs to behave much like their physical counterparts – as reactive objects in my terms -- or they may be programmed to be smart. If they are standard reactive chairs then they will deform

when sat on and move when struck, but otherwise will be inert to changes in their immediate vicinity. In particular, they will be unresponsive to social events such as the number of people trying to sit in them. If they are programmed to be smart – to be active objects -- then they might morph into a couch, or multiply themselves.

2.3   Active objects

Furniture that can self transform, such as tables that can change dimension, lights that automatically move to accommodate extra participants, or bookshelves that shrink or expand, are more complex than reactive objects in two ways. First, they are able to react to changes in social conditions -- the number of people present, the social function of the meeting (e.g. presentation, intimate chat, virtual office meeting, disciplinary discussion, and so on). Second, their reaction goes beyond physical law. They are able to self transform, changing their size, position, appearance, shape and even functionality as appropriate. For instance, if several people enter my office, the books on my bookshelf may transmute into a single icon representing the entire set, the shape of my lamp may alter so that it now projects two distinct beams rather than the one it was projecting before, and my overhead projector, which normally is under the control of whomever is beside it, may change its functionality to now allow non presenters (e.g. participants sitting at the other end of table) to annotate projected slides in whiteboard fashion, provided they are sitting in seats known by the projector to be reserved for participants.

Let us call a virtual object an active object if it is both responsive to interactions with other objects and is able to take actions either on its own or in response to abstract changes in its milieu. Furniture which morphs to meet social needs fits this definition, so do autonomous agents since they are able to autonomously change their position, and orientation, possibly even their internal state, without explicit direction from human agents; and so do human avatars, of course, since they may act unprompted by any apparent changes inside the environment.

Each of these active objects has its own distinctive properties. Autonomous agents, for instance, typically have standing instructions to maintain certain properties of the environment, such as keeping a space tidy, or re-aligning chairs when they are disturbed. They may also have more complex standing instructions to respond to specific requests such as helping humans (or other agents) to locate textual references. (If their role is information related they may be better classified as information objects. But understandably their category membership is ambiguous.) Agents which seem to be reacting to changes in environmental state – e.g. untidiness, or disorder – are not classified as reactive objects because the state they are responding to is abstract, and not governed by laws of physical interaction. For example, an intelligent agent for inventory control that is responsible for restocking its collection of office supplies as soon as its inventory falls below a certain threshold, is responding to a change in the environment, but there are no physical laws about maintaining inventory. Active objects typically have their own agenda and follow their own individual laws programmed into them. Agents which follow or are responsive to social norms also qualify as active objects. Programming such agents is a challenge because the interpretation of social events is a complex matter and typically relies on being able to

attribute motives and beliefs to social participants. Simplified versions of socially aware agents, however, might be able to do such things as move out of the way of an aggressive agent, help another avatar or agent to balance a teacup about to fall, or even stagger back in `horror' when menaced. As virtual environment's become more complex it becomes increasingly important to make agents sensitive to social conditions, though if programmers are not to be behavioral reductionists this will continue to be a formidable task.

For clarity, we shall further distinguish agents which have a regular *physical presence* in a virtual environment from those that are *invisible*, and which are either configured in advance of entering the environment or which can be called up from an environment with their own user interface. Agents which perform tidying tasks such as furniture straightening, or agents configured to maintain the look and feel of a certain room, , are examples of invisible agents which need be called up only when there is a need to modify their agenda. Butler-like helping agents which perform various office duties and which keep a regular physical presence are examples of visible agents.
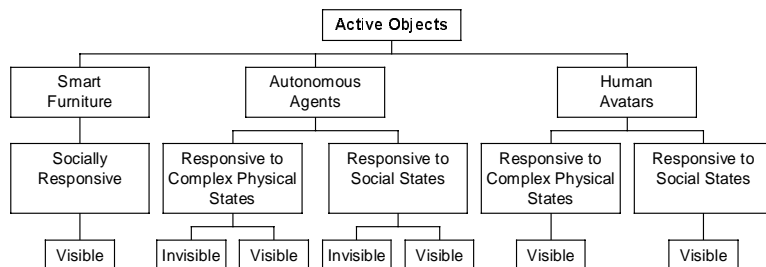

Fig. 1. Here is a simple hierarchy of active objects found in virtual environments. What differentiates most active objects from reactive objects is that they are sensitive to social 'laws' or to highly specific changes in other objects. They therefore are capable of self-transformng actions where they morph, iconify themselves, or perform helpful duties.

2.4 Information objects

In addition to active objects there are also two types of active information entities: *information tools* and *information spaces*. Both have a clear functional purpose: to help users to find information they feel would be relevant to their current activity. An information tool is a visualizing tool that allows its user to explore the contents of a document in a special way. It might be as simple as a tool showing all the paragraphs that contain a given phrase or word, or it may display certain statistical properties of the document. It is a device for visualizing contents. We may imagine it as akin to a lens that can be run across an open document or an unopened book.

An information space allows users to display meta-information about books, articles, magazines, web sites, and to list this meta-information and sort it. In Adaptive Rooms the empty space in a bookshelf, or a corkboard, would likely serve this function. In its most obvious form, the bookshelf space behaves like the display field of a

file system, with tabs that can sort entries by date, size, type, title, and so on. Users can invoke a search of the web, digital libraries and so on, or perhaps autonomous agents will be programmed to retrieve sought after material and post pointers in an appointed information space. The information displayed is active. Just as html documents contain url's to other documents, so the display of meta-data in information spaces are hypertext pointers that, when clicked on, result in the relevant book, article etc. arriving in a physical form and lining up on the shelf. Once the books have arrived on the shelf they can be sorted according to several categories: date, author, title, keyword etc. So the entire bookshelf containing books, articles, and other documents, as well as the open space at the end of the books grouped together serves as an information space where information relevant operations can be performed.

2.5   Does cloning and linking complicate an ontology?

It was mentioned earlier that the physical workflow of an assembly task undertaken in a virtual environment may be different than its physical correlate because it is possible to clone parts and assemblies in virtual worlds, and so distribute problem solving among collaborators in ways that have no counterpart in the physical world. This means that objects in a virtual environment are no longer individuals, but rather types that can be tokened at will. At first this may seem to be no stranger than multiple versions of a document. In collaborative writing, for instance, each collaborator soon develops his or her own version of a document. This has the consequence that there is either a problem of putting together the final version, or there are tight restrictions on who may write to the canonical version. But in assembly tasks, there need be no intermediate canonical versions that must be shared since there may be multiple solution paths. This means that there can be a wild proliferation of partial clones during assembly and no need for version control. But then how are we to characterize such clonable objects? Do they force us to revise our ontology?

The matter becomes even more confusing ontologically, when we consider other tasks, such as jigsaw puzzling. In solving a jigsaw puzzle collaboratively, we may allow players to individually search for placements on their own version of the game, but to enjoy immediate updates of their boards when any of their collaborators find good placements. In simple jigsaw puzzles, it is possible to hill climb to the goal without concern that there may be false or misleading placements. This has the implication that there is no need to assign version numbers to partial puzzle solutions at all: any advance can be shared by everyone. Accordingly, multiple copies of the game can be causally linked, so that a change in one game is propagated to all. But then what sort of ontological status will such multiplicities have?

Such questions raise important issues for designers of collaborative environments. Any theorist eager to explore the ontology of virtual environments must decide how to treat versions and linked multiplicities. For the moment, however, we may note that multiple instantiation, or cloning applies to all virtual objects, regardless of category. Passive and reactive objects may be cloned as well as active objects. Whether human avatars may be cloned is a design choice.

2.6   Interim summary

Assuming that our ontology represents a useful starting point we can say that Adaptive Rooms are collections of virtual objects, mostly self-transforming objects, housed in an architecturally active room with information spaces. The entities of Adaptive Rooms are sensitive to interactive laws of both a physical and a social sort, so many of them respond in a coordinated way as if under the control of a thoughtful butler, always looking for ways to make one's task easier.   It is in that sense that self-transforming objects are the virtual counterpart to smart objects and smart appliances discussed in the literature on Ubiquitous computing. Our objective is to design the objects in our virtual world with enough intelligence and coordination to act politely, to anticipate users' wishes, and to change their behavior with the context of action. Since there is always a danger that too much cleverness may get in the way of effective collaboration, we can only determine the cleverness threshold of smart objects and tools by considering the way they fit into the work flow of collaborators and the resources they use.   In short, the way to design a new collaborative environment is to first understand the way human agents and their everyday objects, representations and gadgets constitute a system of distributed cognition. It is to that we now turn.

3   Work Flow Analyses and Distributed Cognition

To construct useful Adaptive Rooms requires a deep understanding of:

1. How people use resources individually and jointly as they go about their tasks in everyday environments. This is the pattern of work that should be facilitated in an Adaptive Room.
2. How users will adapt the way they do their tasks once they have adaptive virtual rooms to help them. Put differently, we must understand the way external users will interact with Adaptive Rooms as part of their larger ongoing activities.

A few quick examples of the way people use resources around them in everyday environments will help to illustrate why the type of understanding required by point 1. is particularly challenging.

In observing agents interacting with their environments, we have noted that they are constantly managing the resources around them. (Kirsh 96). Much of this activity is quite naturally pragmatic: it is directed at getting things done and accomplishing goals. For instance, people choose, reach for and pick up books and notepads in order to draft memos and outlines. They build models to help solve problems, and they move objects such as chairs around to make it easier to perform their tasks.   The rational structure of pragmatic activity is familiar enough: first complete the sub-goals of a task and then combine these sub-goal states to complete the overarching task. It is reasonable to hope that someday many of the component parts of our pragmatic activities will be automatically carried out by agents in a virtual environment. The idea of smart butlers fits that model. The user initiates an action, perhaps by request, perhaps by gesture, and the butler infers some of the necessary sub-goals and so completes the drudge work part of the activity.   Our research has shown that a surprising

amount of action, however, is not really pragmatically oriented. People seem to spend much of their time performing actions that serve cognitive or epistemic functions rather than pragmatic ones. [Kirsh and Maglio 95]. These actions do not fit into the standard goal sub-goal model of planning and rational action. They are more tightly coupled to the memory, reasoning, and perceptual systems of the people performing them. Hence although some non-pragmatic actions might be anticipated by a smart butler many would lose their effectiveness if executed by someone other than the cognizing agent.

Among the many non-pragmatic activities we have observed and documented include:

− Positioning objects to draw attention to opportunities (e.g., card players arrange their hands to encode plan fragments and possible opportunities).
− Arranging objects to cue action selection or constrain options (e.g., lay out parts in the order they need to be assembled).
− Configuring objects to "remind" users of plans or intentions (e.g., leave keys in shoes, film by the door).
− Preparing the workplace so that the average complexity of tasks that must be done is lower -- amortize the complexity of tasks by paying an up-front cost in preparation -- (librarians pre-sort books to be shelved on carts before moving to the stacks);
− Pointing with hand or mouse to help direct attention or increase the acuity of perception, as in counting coins, or looking for a name in the phone book. (Kirsh 95b);
− Rotating objects physically to save the need to rotate them mentally (rotating tetris pieces (Kirsh and Maglio 95), righting upside down photos (Kirsh, forthcoming));
− Clustering items to create ad hoc categories that are problem-salient (e.g., organizing articles on one's desk to highlight those that are relevant to a current project).
− Re-arranging objects, such as Scrabble tiles, to self-prime or self-cue recall.

In each case what is noteworthy is that agents make short and medium term adjustments to their environments for cognitive reasons: to keep processes manageable, and on track. The conclusion to be drawn is that the environment is one's partner or cognitive ally in the struggle to control activity. Although most of us are unaware of it, we constantly create external scaffolding to simplify our cognitive tasks. Our goals as designers is to create digital environments which make this cognitive alliance as powerful as possible. Helpful workflow analyses must focus on how, when, and why this external scaffolding is created.

3.1   A simple scenario

One collaborative context we are working on focuses on the joint efforts of a marine ecologist, Mr. E, and a geographer, Ms. G. This is a familiar situation: two researchers, each an expert in their own field, join forces to tackle a problem that neither one can quite handle on their own. Ms. G lives on the east coast but regularly visits Mr. E in his office on the west coast. The physical room they typically collaborate in -- E's

room -- has a desktop PC, a Silicon Graphics machine, as well as E's laptop. There is a whiteboard, corkboard, shelves, books, calculator, ruler, three chairs, and a host of sundry other objects, which are in constant use. Because G knows cartography and the geographical information system packages, while E specializes in more theoretical issues, G tends to take the driver's seat in front of the Silicon Graphics machine when they are involved in making a map, with E looking over her shoulder, or working on different material on his desktop PC. E also keeps a small bound notebook with useful parameters tried in the geographical information system package. Since this information helps them keep track of their inquiry, G and E must explicitly resolve their schedules to ensure that this book is where G may find it when she returns to the room and E is not there, or else the must find some other means of guaranteeing that she finds out the parameters E tried when he was on his own. Naturally there is much more to be said about the social arrangements that holds between the two and to the different resources they use individually or jointly. The whiteboard, for instance, plays an important role when E is talking but considerably less so when G is talking.

From our preliminary studies a few non-obvious facts have emerged about the mechanics of collaborative workflow. First, opportunistic encounters in the hallway, outside E's office, and near the coffee machine are important catalysts for their work together. One often thinks of collaborative meetings being scheduled or at least unproblematically occurring. But because much important work occurs during unscheduled conversations, special efforts have to be made in virtual environments to create the right opportunities for informal contact. Second, it is apparent that both E and G take the opportunity to look over the other's work when one of them is out of the room. Data maps are dense and the information represented on them requires a fair bit of study before a viewer feels fluent with them. Accordingly, collaboration is a much larger activity than simple face to face meeting. The state one leaves one's office in, or the state of one's computer screen can be extremely informative to one's partner. Coming early to a meeting can be helpful as means of preparation. But then since helpful state may be stored in notebooks, corkboards, whiteboards and ink marked books, the use of all these resources must be temporally coordinated. Users would like to `page through' the previous stages of work -- not change by change, but collection of changes at a time. Markings that go together -- either by date or content -- ought to be retrieved in a connected way, saving the user from figuring out when a mark on the blackboard was part of one conversation or another, and saving them from having to personally know where to look for relevant documentation (in cases where, for instance, a mark has been placed on whiteboard or notebook after a book has been consulted). Moreover, since these are all resources that may be `messed with' during ordinary non-collaborative activity, it is important to be able to restore or maintain the state of these books despite the other activities one might perform in the same space.

Admittedly, nothing said here is detailed enough about the tasks each participant performs to expose the cognitive or physical workflow of E and G's personal and collaborative activity. But even this cursory account may suffice to show that virtual collaboration may be better than face to face collaboration if it offers participants the capacity to intelligently recover the structure of each other's workspace.

4   Conclusion

Our designs of Adaptive Rooms are guided by three principles:

1.   Wherever possible Adaptive Rooms should adapt to humans rather than humans to them;
2.   Adaptive Rooms should be populated with cooperating smart objects and tools to make our virtual environments like digitally enhanced futuristic ubiquitous computing environments.
3.   An adaptive room and the agents in it constitutes a system of distributed cognition. Problems are solved by coordinating resources in and outside of peoples' heads.

By following these principles we believe that the insights of ubiquitous computing -- that computerization should be everywhere, transparently incorporated -- can be merged with the insights of augmented reality -- that everyday objects can be digitally enhanced to carry more information about their use. The result is an integrated virtual environment in which people can cooperate and collaborate on projects that leverage the computational strengths of computers.

References

1. Bay, S. L. (1979). A study of conditions encountered when traditional educators work in an open space school. Experiential Education. Mankato State University.
2. Kirsh, D. (1997). Interactivity and Interfaces. Instructional Sciences, 25, 79-96.
3. Kirsh, D. (forthcoming). Are all Actions Exploratory or Performatory? Ecological Psychology
4. Kirsh, D. 1996. Adapting the Environment instead of Oneself. Adaptive Behavior, 4:3/4, 415-452.
5. Kirsh, D. (1995a). The Intelligent Use of Space. Artificial Intelligence, 73, 31-68.
6. Kirsh, D. (1995b). Complementary Strategies: Why we use our hands when we think. In: Cognitive Science Society Proceedings. Mahwah, NJ: Lawrence Erlbaum.
7. Kirsh, D., & Maglio, P. (1995). On Distinguishing Epistemic from Pragmatic Actions. Cognitive Science. 18, 513-549.
8. Le Corbusier, & de Pierrefeu, F. (1958). The home of man. Translated by Clive Entwistle and Gordon Holt, London: Architectural Press.
9. Weiser, M. (1993). Hot Topics: Ubiquitous Computing. IEEE Computer, October.
10. Weiser, M. (1993). Some computer science issues in ubiquitous computing. Communications of the ACM, 36,7: 75-84.