A COMPUTER NETWORK MODEL OF HUMAN TRANSACTIVE MEMORY

DANIEL M. WEGNER
University of Virginia

Several of the design factors that must be considered in linking computers together into networks are also relevant to the ways in which individual human memory systems are linked into group memory systems. These factors include directory updating (learning who knows what in the group), information allocation (assigning memory items to group members), and retrieval coordination (planning how to find items in a way that takes advantage of who knows what). When these processes operate effectively in a group, the group's transactive memory is likely to be effective.

We all have known and grown accustomed to the computer model of mind in psychology. It is to the point now that we don't even wince when we hear that our brains are "hardwired" for this, or that our minds have been "programmed" to do that. The scientific advances offered by this model go well beyond these idioms, of course, as the modern understanding of the inner workings of mind appears to be moving ahead rapidly given this rich metaphor. And it's not over yet. There is a whole new level of metaphorical insight that becomes available when we apply the model to the *outer* workings of mind. Quite simply, people may interact with one another in some of the same ways in which computers interact with one another. If individuals are computers, every social group is a computer network.

This article is about how communications among individuals influence the memory processes and capabilities of individuals and groups. The first order of business is an overview, based on the computer metaphor, of the memory systems that are formed when individuals' memory systems are considered together with the communications that

occur between individuals in dyads and in groups. These larger memory systems have been called *transactive memory* systems (Wegner, 1986; Wegner, Erber, & Raymond, 1991; Wegner, Giuliano, & Hertel, 1985), and their influence will then be viewed in terms of issues suggested by the computer network metaphor. The metaphor makes clear that certain elements—what we will call directory updating, information allocation, and retrieval coordination—are necessary for the formation and successful operation of transactive memory. As we shall see, people in groups use many of the same strategems to enhance their transactive memory systems that network engineers have designed for computer communication. With an appreciation of these strategems we can begin to understand how communication influences the memory capabilities of both groups and individuals.

HUMAN GROUPS AS COMPUTER NETWORKS

Normally, using a metaphor involves the application of something you know well to explain something you know less well. In this sense, the metaphor featured here is backward indeed, more like a rohpatem. Most of us, save the occasional computer nerd, are far more familiar with social groups than we are with computer networks, and it seems something of a throwback to try to explain social groups in terms of these networks. For that matter, the computer network literature is itself built largely on the assumption that computers should be linked together in much the same way that people are linked together. So, the backward flow of this metaphor is underway and continues even now as computer networks are advancing in all directions.

So why proceed? The simple answer is that thinking about human groups as though they were computer networks offers us some conceptual distance from our everyday knowledge of groups. Like the exercise of imagining how we humans would appear to a Martian (imagine explaining wax lips or toaster cozies), the project of thinking of ourselves as linked computers offers some realizations that we might otherwise mistake for common sense and take for granted. The more subtle answer to why we would want to develop this metaphor is that in discussing how computer networks are built and the problems that must be overcome in engineering them, we may learn something of the obstacles naturally encountered (and solved) in the evolution of human group memory systems.

Much can be learned about a system by attempting to build one. This is the basic idea underlying the use of artificial intelligence techniques for the construction of psychological theory, and the notion of theorizing by engineering is now well-developed in the study of individual cogni-
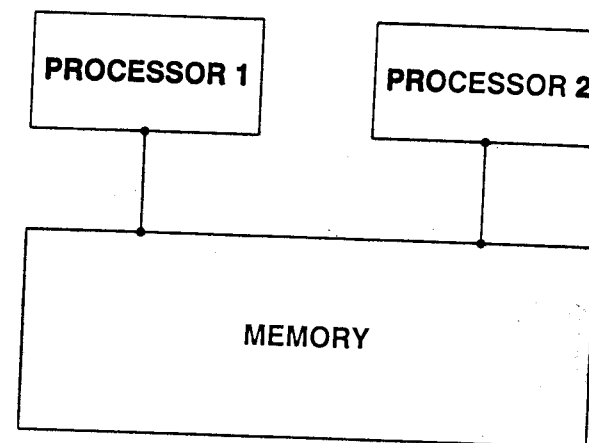
FIGURE 1. Distributed processing with memory sharing.

tive process (e.g., Braitenberg, 1984; Minsky, 1986). This approach has not been applied with much regularity to social systems (but see Latané & Nowak, 1994; McPhail, Powers, & Tucker, 1992), or at all to the phenomenon of social memory, and it is about time. So get out your tools, cables, computers, and instruction manual. In the next few pages, we're going to learn about social memory by building one. In so doing, we will construct a new basis for reviewing and extending the analysis of transactive memory.

The design of computer networks is quite sophisticated these days, but for our purposes we only need to explore some basic ideas (see Mason, 1987). To start, let us examine a simple computer network that enjoys the benefits of shared memory. Figure 1 shows two processors that share one physical memory. This might happen, say, in your bank or supermarket, where more than one computer processor can read from and write to the same physical memory or database. The benefit of this system is that each processor has access to information that is perfectly current among all processors; every clerk in the bank will get the same results when checking into your balance, and every supermarket checker will put the information in the same place when acknowledging that you have paid for your bananas. The manager can look into memory to see how much of what is being sold and so can order more without even looking at the shelves. The only bottleneck in this arrangement is the speed with which an individual processor can write to or read from memory. And the one minor problem that must be solved along the way is preventing two processors from acting on the same memory item simultaneously. All
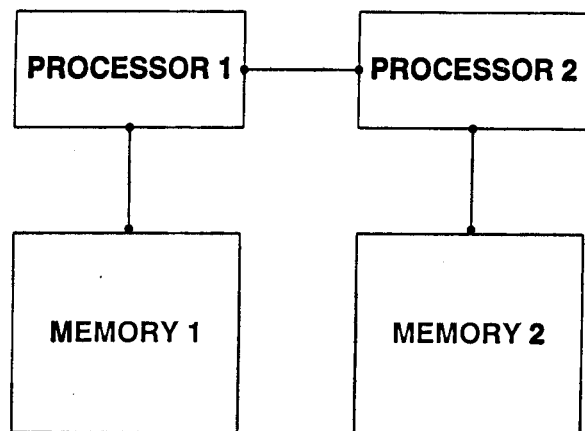
FIGURE 2. Distributed processing with message passing.

that is needed is some "under revision" flag that each processor puts out at a memory location to keep the other(s) from accessing or changing that location while the first processor is doing its work.

Although it seems exceptionally useful to have this sort of mass memory bank, sadly this system corresponds to no human system currently known. People don't have physically connected memories, nor do they possess the ability to read or write directly to each other's memories. Social theorists have long imagined this possibility in their conceptions of the "group mind," however, as it is very useful to represent group behavior as resulting from some common bank of information (Wegner, 1986; Wegner et al., 1985). Unfortunately, this kind of thinking tends to inspire overenthusiastic extensions that propose a real group memory, rendered if not in protoplasm then in ectoplasm (e.g., Jung, 1922; Pareto, 1935), and such extensions clearly miss the point. We are not all joined at the head. The way social memory actually looks is modelled better by the computer network shown in Figure 2. This network involves message passing between processors, each of which has its own physical memory.

Like a computer that sits on the desk at work, connected by modem or network cable to another computer elsewhere, this network offers a form of connectedness that has some benefits. To some degree, the two memories are in contact. One processor can reach information in the other processor's memory by asking the other processor. (This system is also a bit like what happens when you have computers at home and at work that you "connect" by carrying a floppy back and forth every day to transfer files from one to the other.) The rudimentary benefit to be gained here is that there is at least basic communication available between memories.

However, there are three serious disadvantages to this system as it currently stands. First, there is a new bottleneck that is likely to make everything painfully slow. Whenever Processor 1 wants something in Memory 2, it must first request it from Processor 2 by passing a message through the potentially sluggish modem line (or worse, the hand-carried floppy), and then must download it from Processor 2—which involves another transfer. The second disadvantage is that this operation ties up both processors. Unless the computers can multitask (do multiple things at once—which humans also may or may not be able to do, depending on context), they will both be completely occupied by the mere task of writing to or reading something from social memory.

The third disadvantage of this system is the most disabling. With only the processor connection, the interaction between computers is blind. Neither processor has a clue about whether the other's memory will contain any given item, and so many messages will be either foregone or sent for no reason. Processor 1 might be doing something that requires knowledge of the price of noodles, for example, and this price might be in Memory 2. But Processor 1 would not know this, and so would fail in its noodle-relevant task. Each processor might thus never take advantage of the potentially useful memory banks available through the other. A processor might also be programmed to seek out every unknown item from the other processor. In this case, the processor working on a task might ask endless questions of the other, to the point that the second processor would be totally devoted to the same task occupying the first—the equivalent of an afternoon with a six-year-old who has just figured out about "why?" questions. The only way a system like this can work is if a wise human is running the whole thing and spends time keeping track of what is in which computer.

In some computer networks, a solution has been developed for this problem that begins to make the system work as though all the processors had physical access to one big memory (as in Figure 1), even though they are physically separated. The network plan shown in Figure 3—message passing with directory sharing—allows individual computers to be connected so as to benefit from a virtual shared memory. In essence, it uses software and a processor connection to simulate a massive hardware memory. This system functions on the principle that any computer's memory is not usually just a random jumble of information, but instead is organized in terms of some set of categories or files. The list of these files is a directory, and as a rule the processor reaches its own
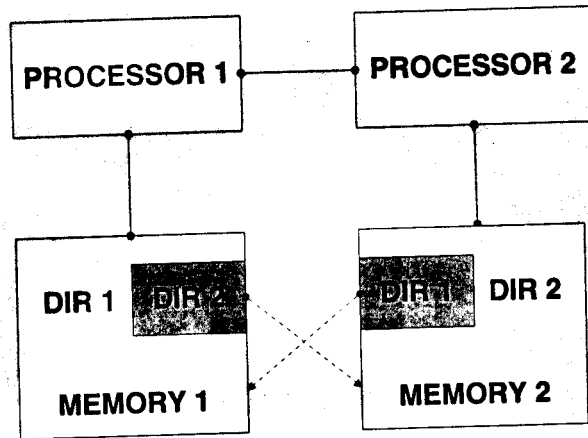
FIGURE 3. Distributed processing with message passing and directory sharing.

individual files or items by addressing them in their appropriate directory entry. All that is needed to produce a slick networked memory is to have the directory representation of a computer's memory contents that is available for the memory's own processor also made available in some form for the other processor.

As shown in Figure 3, each memory in this network contains its own directory, and also contains a directory of the other memory. When Processor 1 then queries its memory for a file, for example, it can review not only whether this file will be available in Memory 1, but also whether it will be available in Memory 2. If the file is in Directory 1, the processor can find the information immediately in its own memory. If the file is not in Directory 1, but exists in Directory 2, Processor 1 can then ask Processor 2 to query Memory 2 for the information. Although a processor's retrieval of information from distant memories is substantially slower than retrieval from its own, the quick searches of own memory can be supplemented when needed by these additional searches, to yield a far larger potential database.

There are three issues that must be dealt with if this more elegant network is to work. First, there is the issue of *directory updating*. Updating the directory to a processor's own memory is not a problem. Whenever items are stored or retrieved in the processor's own memory, they are addressed by means of the directory and are thus intrinsically linked to the directory. The difficulty comes in updating the directories to the memories linked to other processors—without actually storing, retriev-

ing, or reviewing all those memory items individually. One solution to this problem is to make each processor review the other's directory entries by virtue of message passing whenever both are not busy. Another is to arrange that the directory structures of both processors are identical, such that when one accesses its own memory through a directory and finds no relevant item it immediately turns to the parallel directory in the other machine. Even this solution requires directory updating, however, because any new directory structure occurring in one machine must be duplicated for the other. Regardless of the particular solution, then, directory updating is a key activity that improves both the speed and breadth of information availability in a network of this kind.

The second issue of importance for this directory-sharing network is *information allocation*. What happens when an item arrives in the network? It could be that each processor that receives input of any kind would just dump that information immediately to its own memory, as this procedure would be more efficient at this point than any attempt to allocate information to another memory. And, if the two directory structures are indeed identical, as is one option in this network, then it matters little where any particular item is stored. The first processor to get it should just keep it in memory. However, if the directories diverge in any way, the long-term efficiency of the network would be improved if, in periods when both processors are free, they "clean house" by transferring information from the computer with a less sophisticated directory structure for that information to the computer with the more sophisticated structure. If one machine has everything on its hard disk in one big directory, for example, and the other has a specific directory entry for noodle information, the system as a whole would be able to find noodle facts more quickly if any noodle items were transferred to the noodle directory in the machine that "knows noodles"—and a notation of that directory's existence were entered in the other machine(s). Such allocation of information also solves in advance any problems that might arise later from disagreements between multiple sources of information in the network; multiple files on a topic could differ because some are not current or are early drafts, and a single file in a single location avoids this ambiguity.

The third issue that needs attention in the operation of this system is *retrieval coordination*. A system like this, once built, is only accessed efficiently if each machine is programmed to follow retrieval rules that take advantage of the system. If one machine needs a memory item that it cannot access internally, it needs a plan for deciding when to query the other computer; and if there are many other computers, it needs a further set of decision rules for selecting which to query first, second, and so on.

It is often a matter of judgment whether an item will be found under a particular directory label, and, when there are multiple possible directories, even a fully updated directory system may not make it obvious where the item might be found. In this sense, each computer needs a coordinated internal model of the directory structure of the whole system. The alternative is a random search that will tie up all the processors and undermine the system's efficiency.

By now, we're far enough into computer network theory that any reader who has not wandered off into traffic must already understand the point: Human beings in pairs and groups form message-passing, directory-sharing memory systems. This means that in social networks, we get many of the benefits of having greatly expanded memory systems, while suffering little of the reduced speed of retrieval that might befall us if these systems were only built on disorganized chatter.

The basic insights we can take away from the computer network metaphor for understanding human transactive memory correspond to the three issues noted above for the architecture of message-passing, directory-sharing systems. First, the formation of human transactive memory systems depends on processes whereby people learn what others are likely to know about (directory updating). Second, the formation of these systems also requires processes whereby information coming into a group is communicated to individuals whose expertise is likely to facilitate its storage (information allocation). Third, the operation of human transactive memory depends on each person in the group having a retrieval plan for any topic based on the relative expertise of self and the others in the group such that the sequence of locations for any information search can be determined (retrieval coordination). In short, a complete transactive memory in a group occurs when each member keeps current on who knows what, passes information on a topic to the group's expert on the topic, and develops a relative sense of who is expert on what among all group members. We now discuss each of these processes in turn.

DIRECTORY UPDATING

Our directories for memories held by others can be thought of as metamemories (Flavell & Wellman, 1977; Nelson & Narens, 1990). That is, they are memories about memories. But they are not memories or judgments about our own memories—they are memories about the memories of others. Just as we can separate knowing something from knowing that we know it (Hart, 1967), we may know that others know something without knowing it ourselves. The question then becomes: How do we know they know it?

People rely on several different sources of information in creating and then updating their transactive memory directories for any group (cf. Wegner et al., 1991). At the most basic level, a directory for an individual could be created with *default entries* based on the observation of surface characteristics of the individual. When we are in an impromptu gaggle waiting for a bus, for example, we can probably can make some rough estimate of who will know about automobile tires (the men rather than the women?), about basketball (the tall rather than the short?), or about heating pads (the old rather than the young?) even though these guesses are based on stereotypical overgeneralizations and are bound to be wrong and potentially prejudicial in many individual instances. When there do exist differences in an area of memory expertise identifiable through physical characteristics, however—for example, in that women often know more than men about their social relationships (Ross & Holmberg, 1988)—it becomes useful for individuals in groups whose members are otherwise unacquainted to base their searches for memory items on observations of such surface features.

One of the key purposes of conversations in working groups is the updating of individuals' directories for group information beyond this default level. Oftentimes, this means making explicit assignments of directories to individuals ("Norton, you count the mops"). Individuals may also volunteer or otherwise accept responsibility for a range of information. In all these cases it can be said that *negotiated entries* have been created in the directories of all who hear of this allocation. When these assignments are public, they become immediately useful to all who know of them. This kind of planned knowledge allocation seems particularly likely in task groups, but it can also happen when groups formed for any purpose find they can perform that function more effectively by planning explicitly for differential expertise. Couples in close relationships may negotiate who will be the financial expert and do the bills (Atkinson & Huston, 1984), for instance, or who will be the Halloween costume expert and remember where these are when they are needed.

It is quite unlikely, though, that explicit negotiation or conversation will be directed toward many such items. The "Halloween stuff job" just falls to someone in the family, and it would be strange to talk about it ("Dear, I'm so very tired of keeping track of the pointed hats"). More often, the division of whole domains of information happens on the basis of some perception of differential expertise or interest among group members. These *expertise entries* mark one person as having more knowledge or interest in an area than do the others, and this person becomes the group's designated storage location for items in that area. In groups with clear role distinctions, the roles often exist as markers of expertise,

but in unstructured groups, the attribution of expertise can occur more capriciously. It doesn't take much to get this rolling, because the perception that a person knows about one item in a category may often be enough to get the ascription of expertise underway. Just as a person with one cow figurine becomes the owner over time of a cattle trinket museum—because everyone thinks the person is a cow enthusiast and so brings cattle artifacts from distant airport gift shops—the person in a group who knows one thing about cattle may become the group's cow expert.

Directory updates also can come from knowledge of the access to information that others enjoy. Such *access entries* occur when a person's exposure to information is known to the other. Knowing that a person has accessed information earlier than have others is one such indication. For example, such *primacy* is operating when the first person in a group to hear about the weather prediction for the weekend is expected to keep track of the weather updates later on. *Duration* of access is also a factor, such that the person who is in Wisconsin longer than the others, for example, is more likely to be expected to know where to find Beloit. And *recency* of access can also be a cue to knowledge. The person who most recently had the TV remote control is expected to be able to locate it even though others may have had it first or for much longer. Wegner and associates (1991) note that these rules seem to operate for citations in scientific journals. When one is hoping to make a citation for a particular point in an article, the usual procedure is to cite either the researcher who found this first, the one who has devoted an entire career to it, or the one who has made the point most recently. In a major review article, one might even cite all three. It is bad form to cite a list of references that meet none of these criteria.

In many groups, these sources of directory entries are all that are required for the formation of effective directories for each member that provides a map to information in the group. When the information to be stored in a group becomes particularly fine-grained or poorly organized, however, it may be that some individuals will be needed to become directory specialists for the group as a whole. The receptionist or operator who works the phones in an organization, for example, may need to keep a far more detailed directory of the areas of knowledge held by group members than will any other member of the group. When such specialization occurs, the ability of this person or persons to keep up with the changing face of individuals' memories may require extensive directory-updating activities. The recent proliferation of touch-tone telephone menu systems ("Touch 1 to speak to your Mom") accentuates the difficulty of maintaining such directory expertise in enough human receptionists to allow many people easy access to the information held

in any large organization. The directory specialist thus has a certain form of power in the group that should be obvious to anyone attempting to access the group's knowledge (see Krackhardt, 1990).

A final complication here is that the updating of directories is not always easy. In a couple who have been together for some time, for example, each partner will have formed a fairly stable directory to the other's memory storage. Quick changes in the couple's scheme are likely to be difficult to incorporate, as the grooves of habit will be deep. When he has always taken care of the rabbits and she has always tended the mule, both rabbits and mule might go hungry for a while if the pair suddenly switched chores. Wegner and associates (1991) found in this regard that while offering a new memory organization scheme—an imposed directory—to pairs of unaquainted people increased their group memory ability for items that fell into the scheme, offering such a scheme to a close couple not only didn't help but actually hurt their memory performance. If directory updates are to have the intended effect of enhancing group memory, they must occur a little at a time such that the group memory structure can itself be remembered.

INFORMATION ALLOCATION

It is risky to pass information from one person to another. The classic studies of Bartlett (1932) revealed that information is quickly degraded when it moves along a chain of people, often to the point that it is no better than information that has been moldering in one person's memory for many years. This disadvantage suggests that it might be most efficient for groups if the first person to receive any item of information is generally assigned to keep it in personal memory. However, this tactic undermines the formation of effective group memory structures because it allows memory to become scattered among group members. Each person's directory becomes very large and haphazard, and so becomes difficult for others to remember. For this reason, the risk of information degradation is often accepted as a tradeoff for the benefit of efficient group memory structure. Information is passed to the person in the group who has the most relevant and most well-developed directory structure for items of that kind.

Ideally, the information is passed immediately, perhaps even without being encoded in any other memory. The new cookbook is rapidly conveyed to the family cook, for example, or the bills that arrive in the mail are placed on the family money-manager's desk. Any delay in getting the item to the expert means that it will not be accessible to other group members. In these instances, the person providing temporary storage of the information may know it only by some label while it is in

transit (e.g., "the cookbook"), and so will not have contact with the lower-level items of information (i.e., specific recipes) at all. Only when the items are actually in the memory of one of the group members (i.e., the cook reads the book) will the group have use of them (i.e., fewer poptart-based meals). The kind of information passing that occurs in this way—without initial storage of the lower-level information on the part of the passer—is especially efficient because it involves passing information without the passer necessarily memorizing the information personally or subsequently forgetting it personally.

There are less efficient but more interesting forms of information passing that involve seemingly complete transfers of lower-level information. One person may be holding some information in memory, only to learn that the other person knows much more about this sort of thing. The first person passes along the information in detail to the second, and then proceeds to ignore these matters from then on. This might happen when, for example, a researcher who has worked at a small college for many years and has always devised his own grant proposal budgets gets a job at a large university where the department has a budget specialist. From then forward, the researcher need never worry about calculating fringe benefits again, or about dealing with all the other minutia that must be reinvented with each proposal submission. One result of this is that the budget specialist will learn a bit more about how to do budgets from this newcomer. The other result, however, may be that the researcher soon forgets budget-writing entirely. Items that are not regularly retrieved from individual memory, after all, may become less retrievable no matter what their initial strength of encoding (Bjork & Bjork, in press), so the person who has "passed" the information may actually lose the ability to retrieve it over time. In this way, transactive memory may have the effect of creating remembering in some individuals and forgetting in others.

The eventual consequence of these processes of attention and disattention and remembering and forgetting is that individuals in groups become progressively more specialized for memory tasks. Although everyone may start out equal, with similar directory structures and storage patterns, small disparities in initial directory entries between individuals can become magnified over time by virtue of the information allocation process. This *progressive differentiation* of transactive memory can magnify default directory entries, for instance, to the point that stereotypical memory topics in fact do become exclusively the aegis of the persons for whom they are socially expected. Sex-role distinctions in memory tasks may become exaggerated in relationships for this reason (he *always* remembers where the tools are, she *always* knows what the baby should be fed). At the same time, negotiated entries, expertise

entries, or access entries are also likely to be amplified through this process, such that any minor initial deviations from a "vanilla" memory structure will over time become major as the result of group membership. We often think of eccentricity as a property of someone who spends too much time alone—but here we find a social process that constructs eccentrics of sorts in service of enhancing the memory of the group to which these people belong.

In a recent study of the effects of group training on group memory, Liang, Moreland, and Argote (in press) have observed both the occurrence and the utility of progressive differentiation. Subjects in this research were trained to assemble radios either in work groups or individually, and then later were tested on this task either with their original group or, for those trained individually, in newly formed groups. Groups whose members were trained together recalled more about the assembly procedure and produced better quality radios than did groups whose members were trained separately. Extensive coding of videotapes of the assembly sessions showed, in addition, that group training increased the differentiation of knowledge in the group—in that it increased specialization in remembering distinct aspects of assembling the radio, increased the coordination of the group in their work, and increased how much group members trusted one another's knowledge about assembling the radio. Analyses showed, moreover, that these factors played a significant part in improving group performance in the group training versus separate training conditions.

These findings have been amplified in further research by Moreland, Wingert, and Argote (cited by Moreland, 1994). In this study, in addition to groups providing a direct replication of the Liang and associates' (in press) experiment, two further conditions were explored. In one, subjects were trained individually and then had a brief team-building exercise (and were then tested in these groups); in the other, subjects were trained together in a groups but were then tested later in different groups. It was found that groups in these two new conditions performed as poorly as groups in the individual training condition. Groups in the group training condition did much better, as was found by Liang et al. These control conditions lend further credence to the idea that the development of the group's transactive memory was a key feature in the effect of group training on subsequent performance by the group.

If these results characterize the typical operation of working groups, they suggest that progressive differentiation is a natural and useful aspect of transactive memory formation. Groups that encounter information together become proficient as they learn to allocate the information and then retrieve it when they perform. There are obviously exceptions to this rule. At the extreme, complete differentiation would

be counterproductive. Some knowledge must be shared—as when, for example, everyone in a group needs some item of information so that all can behave in some uniform way (e.g., attend a party). It is also the case that these sorts of differentiations would not be likely to occur were it not for the existence of shared information about the directories of others. If the female in a couple didn't know that the male knew how to fix a leaky faucet, for example, she might need to attend to such information herself. Without the reliance on others for memory that comes with personal directories to their expertise, one needs to maintain one's own expertise and so will not be inclined to specialize.

Differentiation may also be unlikely when the information is of greater interest to individuals than it is to the group. When, for example, a couple discusses the odd fact they just learned about their mutual acquaintance, they may go on for some time comparing facts and filling in the story from both of their perspectives so as to form a representation of this person that is satisfying to each partner (Ruscher & Hammer, 1994). Because the couple has little joint interest in this person, however, the information may not be entrusted to either partner alone and instead may become part of the couple's shared stock of knowledge. Duplicate, undifferentiated memory structures can often occur when there is no special pressure for the group to form an efficient system for storing and accessing information.

One other exception to the differentiation rule involves memory information that is simply not shareable. The examples we have used to this point suggest that everything can be communicated, but this is not so. When an individual learns skills, for example, or acquires information implicitly (without knowing that he or she has acquired it; see, e.g., Graf & Masson, 1993), these items cannot be transferred readily to another person. In the case of skills, time and practice are required for transfer, and even then the skills are not shareable with others in the group in the same way that specific answers to questions can be shared (Argote, 1993). In the case of implicit memories, the individual does not have verbal access to the fact that the learned information is in memory, and the information shows its presence only by influencing performance. Only when items in memory take the form of explicit, verbally-represented memories can they be passed among group members and accumulated in particular persons. These, then, are the memory items that are likely to show progressive differentiation as the result of the information allocation process.

The information allocation process, in sum, is the procedure whereby individual memories are fashioned into a differentiated group memory that is useful to the group. Individuals are changed in this process, such that they come to know some new things and lose contact with some old

things. This only happens, however, to the degree that the group knows about it. The group must have access to this differentiated structure in the directories of its individuals if the process of information allocation is at all to be useful. Without directories to this information, group members would be in the same position as members of newly-formed groups. They would tend only to talk about shared knowledge, rather than accessing the unshared knowledge held by individuals (Stasser, Taylor, & Hanna, 1989). Only in groups that have been together long enough to form useful directories can it be expected that unshared information will be accessed to the benefit of the group as a whole.

RETRIEVAL COORDINATION

The individual in a group has at least two and sometimes many more directories to consider in deciding where to look for any memory item. The question of retrieval coordination is the issue of how to organize the search process during retrieval so as to maximize both the speed of the search and its likelihood of finding the needed information. There must be an overarching "directory of directories," or at least its functional equivalent, if the most efficient searches are to be made. A large part of this coordination process involves a first decision of whether to look beyond the self at all. Retrieval coordination may invoke, at least initially, then, an assessment of one's "feeling of knowing" (Hart, 1967) for the desired information.

According to accounts of the "feeling of knowing," a question directed to a person can prompt retrieval of the relevant memory item, and can also prompt retrieval of an assessment of whether that item is known (e.g., Nelson & Narens, 1990). The "knowing" judgment is predictive of whether the item will indeed be retrieved, but not perfectly so. Indeed, assessments of "knowing" can be influenced by a variety of variables that are irrelevant to the likelihood of retrieval. But such "knowing" judgments may nonetheless influence whether the individual will trust his or her own item retrieval or instead call on someone else for the item. Although research has not addressed this directly, it would seem that individuals should also be able to generate "knowing" judgments for others in their group—based on their directories for those persons. With a strong feeling that one's spouse would know where the camera is, for example, why should one even think about it where it is before asking?

This reasoning suggests that individuals will often look outside themselves for items of information, sometimes even before they attempt to retrieve the items from their own memory. When the items are simple or brief, of course, retrieval could be so fast and automatic that it overrides any attempt to look to others. When items are complex,

lengthy, or difficult to retrieve from one's own memory, however, one may often turn to others who seem more likely to know. Indeed, there are even some instances in which the first thing that comes to mind regarding some topic is where to find out about it rather than any lower-level information on what it is. When someone mentions a particularly arcane statistical technique, for example, one may think right off of Professor Brontosaurus—the person who taught this technique in a class—rather than thinking of how one might go about doing it. When retrieval questions prompt thoughts of "who would know this?" one can be fairly certain that one's own directory inspires less confidence than the directory of this other person and retrieval attempts will be interpersonal rather than intrapersonal.

Effects consistent with this proposition have been observed in several studies. Smith and Ellsworth (1987), for example, found that people were more inclined to change their memory reports in the direction of misleading questions when those questions were asked by an examiner who was ostensibly knowledgeable about the events being remembered. Loftus, Levidow, and Duensing (1992) surveyed visitors to a science museum, in turn, and found that those who are inaccurate in a particular memory domain also tend to be suggestible in that area. Children and the elderly, for example, were relatively inaccurate and also suggestible. Artists and architects, on the other hand, were suggestible despite being relatively accurate. The assessment of one's own relative memory abilities need not reflect reality completely—although such misassessments would tend to make one less useful as a member of a transactive memory group. Asking others for items one already knows, or ignoring their retrievals when these are more correct than one's own, are good ways to undermine group memory performance.

The retrieval coordination involved in deciding whether self or other is a better source of memory may underlie the finding that group memory is generally better than individual memory. In an early study of this, Yuker (1955) had individuals recall their own versions of Bartlett's (1932) story "The War of the Ghosts," and then had them discuss this in small groups and generate group recall protocols as well. The group versions were more accurate than the average individual recall in 38 of 40 groups tested. This finding has since been observed across research in this area (Clark & Stephenson, 1989; Hartwick, Sheppard, & Davis, 1982; Hinsz, 1990; Vollrath et al., 1989). If individuals follow the strategy of withholding their answer whenever their own feeling of knowing for an item is weak, any answers provided by other group members might be more likely to be correct and would be substituted to make the group performance better.

The fact that individuals' feeling of knowing judgments are not per-

fectly correlated with their actual memory would suggest, further, that group members would not always be inclined to dismiss their own (inaccurate) retrievals in favor of those of others. This means that while there is some noteworthy improvement from individual to group memory, this improvement is not complete. That is, the group memory will not necessarily capture all the accurate information that individuals have in their memories to create a "best possible" retrieval. Research by Hinsz (1990) and by Vollrath, Sheppard, Hinsz, and Davis (1989) indicates that newly-formed groups indeed do not perform up to this "truth wins" level.

A key aspect of retrieval coordination, then, is the development of the recognition of expertise among group members. Libby, Trotman, and Zimmer (1987) found that work groups could recognize individual members' expertise to some extent, in that the group member identified as having the most expertise at the task did perform significantly better than the average member. It also appears that the recognition of a group's expertise overall is augmented by group size. Littlepage and Silbiger (1992) found that larger groups were more likely to recognize their own expertise on a task.

Another finding in the group memory literature is that groups inspire individuals to report greater confidence in memory (e.g., Hinsz, 1990). This makes sense, too, as a concomitant of the retrieval coordination process. The individual in a group achieves a feeling of knowing not only for what he or she knows, but also for what other group members know. This group-based feeling of knowing may color an individual's estimates of knowing, especially if the individual has been a member of a group for a long time and has become accustomed to retrieving information in the group transactively. When this happens, the perception of boundaries between one's own directory and the directories of other group members may be undermined. It may even be that this process is involved when members of a group come to appropriate the ideas of other members (cf. Wicklund, 1989).

These kinds of observations remind us that retrieval coordination is an individual enterprise, something that happens inside the head of each group member, and that therefore can be more or less effective given the individual's proclivities. Individuals might claim more expertise in certain areas than others for egotistical reasons (Greenwald, 1980), they may take different perspectives on social events that influence their memories and their appreciation of others' memories (Baumeister, Stillwell, & Wotman, 1990), and they may contrive to present their memories to others as more or less veridical for social reasons (Gentry & Herrmann, 1990). These forces are likely also to sway the individual's efforts at the coordination of directories in a group, such that the person's desired or

preferred organization of knowledge in the group may influence the search for information more than the actual organization of such knowledge.

Retrieval coordination may break down to the degree that the individual turns to others for retrieval of information which the individual himself or herself is in a far better position to provide. The tendency of children to accept memory suggestions from family or from their interrogators is now well known (Ceci & Bruck, 1993), for example, and the developing controversies surrounding the veracity of memories of childhood abuse uncovered years later in therapy (e.g., Benedek & Schetky, 1989) indicate that there are points when people may indeed draw on others' memory suggestions in the process of producing their own recollections. Moreover, since the sources of a memory are not necessarily remembered piggyback with the memories themselves (Johnson, Hashtroudi, & Lindsay, 1993), the potential for confusion of internal and external sources of memories is very real. The intrusion of others' memories or suggestions into one's own memory reports seems especially likely to occur in the climate of trust in the memories of others that exists when own and others' memories combine in a transactive memory system.

CONCLUSION

The central idea of this article is that the interaction of human minds forms a system of memory that has certain important resemblances to a computer network. When we imagine what kinds of considerations must enter into building a computer net, we find several useful inspirations for understanding how this kind of process ensues when human groups form and develop cooperative memory capacities. The computer model shows us that getting organized (directory updating), channeling information to the right places (information allocation), and having a strategy for getting it back (retrieval coordination) are all things we must think about to link computer memories. And as we have seen, there is evidence that these same tasks arise and must be undertaken when humans remember in groups. To the degree that human groups do solve these problems, it appears that their group memory structures develop and become capable of memory feats far beyond those that might be accomplished by any individual.

The computer network model does not capture all that goes on in the affairs of human social memory, of course, as computers themselves are still no more than stick figures against the classic art of the human mind. We cannot learn too much from this model, for example, about affective influences on group memory, or about ways in which human desires or the clash of egos might prompt new forms of remembering and forgetting in groups. The computer network model also falls short in the sense that it offers only possibilities about human interconnections rather than realities. There are many possible ways to simulate a network process, just as there are ways to simulate any mental function, and there is no guarantee that we have selected exactly the same architecture that occurs in human networks. At best, we can appreciate that some of the basic links between our computers, and the functions that these links serve, must have counterparts of some kind in the way humans coalesce their individual memory systems.

It is worth noting, in conclusion, that this article itself is an example of the operation of a transactive memory. Having written on this topic off and on for several years since participating in its development with Giuliano and Hertel, I have become a collector of sorts for any bits of knowledge dealing with the topic of transactive memory. Psychologists who have heard of this idea send me things from time to time or mention them to me at conferences. They allocate information to me in the form of literary quotes, interesting examples from their lives, research findings from their studies, and things they have read that relate to the topic. Without asking for this, I have become a repository for psychology's knowledge of transactive memory, a node on the network as it were. Writing this article, then, is my way of saving people the time it would take to ask me about these things. Please update your directories in response to this article. For the next while—until I have accumulated enough material for a new paper—this *Social Cognition* issue knows everything I know about transactive memory and can take my place as a speedbump on the information superhighway.

REFERENCES

Argote, L. (1993). Group and individual learning curves: Individual, system, and environmental components. *British Journal of Social Psychology, 32*, 31-51.

Atkinson, J., & Huston, T. (1984). Sex role orientation and division of labor early in marriage. *Journal of Personality and Social Psychology, 46*, 330-345.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* Cambridge, England: Cambridge University Press.

Baumeister, R. F., Stillwell, A., & Wotman, S. R. (1990). Victim and perpetrator accounts of interpersonal conflict: Autobiographical narratives about anger. *Journal of Personality and Social Psychology, 59*, 994-1005.

Benedek, E. P., & Schetky, D. H. (1989). Problems in validating allegations of sexual abuse. Part 1: Factors affecting perception and recall of events. In S. Chess, A. Thomas, & M. E. Hertzig (Eds.), *Annual progress in child psychiatry and child development* (pp. 627-636). New York: Brunner/Mazel.

Bjork, R. A., & Bjork, E. L. (in press). A new theory of disuse and an old theory of stimulus fluctuation. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *From learning*

processes to cognitive processes: Essays in honor of William K. Estes (Vol. 2). Hillsdale, NJ: Erlbaum.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Ceci, S. J., & Bruck, M. (1993). Suggestibility of the child witness: A historical review and synthesis. *Psychological Bulletin, 113*, 403-439.

Clark, N. K., & Stephenson, G. M. (1989). Group remembering. In P. Paulus (Ed.), *Psychology of group influence* (2nd ed., pp. 357-391). Hillsdale, NJ: Erlbaum.

Flavell, J. H., & Wellman, H. M. (1977). Metamemory. In R. V. Kail & J. W. Hagen (Eds.), *Perspectives on the development of memory and cognition* (pp. 3-33). Hillsdale, NJ: Erlbaum.

Gentry, M., & Herrmann, D. J. (1990). Memory contrivances in everyday life. *Personality and Social Psychology Bulletin, 16*, 241-253.

Graf, P., & Masson, M. E. J. (Eds.) (1993). *Implicit memory: New directions in cognition, development, and neuropsychology*. Hillsdale, NJ: Erlbaum.

Greenwald, A. G. (1980). The totalitarian ego: Fabrication and revision of personal history. *American Psychologist, 35*, 603-618.

Hart, J. T. (1967). Memory and the memory monitoring process. *Journal of Verbal Learning and Verbal Behavior, 6*, 685-691.

Hartwick, J., Sheppard, B. H., & Davis, J. H. (1982). Group remembering: Research and implications. In R. A. Guzzo (Ed.) *Improving group decision making in organizations* (pp. 41-72). San Diego, CA: Academic Press.

Hinsz, V. B. (1990). Cognitive and consensus processes in group recognition memory performance. *Journal of Personality and Social Psychology, 59*, 705-718.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3-28.

Jung, C. G. (1922). *Collected papers on analytical psychology* (2nd ed.). London: Balliere, Tindall, & Cox.

Krackhardt, D. (1990). Assessing the political landscape: Structure, cognition, and power in organizations. *Administrative Science Quarterly, 35*, 342-369.

Latané, B., & Nowak, A. (1994). Attitudes as catastrophes: From dimensions to categories with increasing involvement. In R. R. Vallacher & A. Nowak (Eds.), *Dynamical systems in social psychology* (pp. 219-249). San Diego, CA: Academic Press.

Liang, D. W., Moreland, R., & Argote, L. (in press). Group versus individual training and group performance: The mediating role of transactive memory. *Personality and Social Psychology Bulletin*.

Libby, R., Trotman, K. T., & Zimmer, I. (1987). Member variation, recognition of expertise, and group performance. *Journal of Applied Psychology, 72*, 81-87.

Littlepage, G. E., & Silbiger, H. (1992). Recognition of expertise in decision-making groups. *Small Group Research, 23*, 344-355.

Loftus, E. F., Levidow, B., & Duensing, S. (1992). Who remembers best? Individual differences in memory for events that occurred in a science museum. *Applied Cognitive Psychology, 6*, 93-107.

Mason, W. A. (1987, November). Distributed processing: The state of the art. *Byte*, 291-297.

McPhail, C., Powers, W. T., & Tucker, C. W. (1992). Simulating individual and collective action in temporary gatherings. *Social Science Computer Review, 10*, 1-28.

Minsky, M. L. (1986). *The society of mind*. New York: Simon & Schuster.

Moreland, R. L. (1994, July). *Training people to work together in groups*. Invited paper presented at the meeting of the American Psychological Society, Washington, DC.

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *The Psychology of Learning and Motivation, 26*, 125-141.

Pareto, V. (1935). *The mind and society*. New York: Harcourt-Brace.

Ross, M., & Holmberg, D. (1988). Recounting the past: Gender differences in the recall of events in the history of a close relationship. In J. M. Olsen & M. P. Zanna (Eds.), *Self-inference processes: The Ontario Symposium* (Vol. 6, pp. 135-152). Hillsdale, NJ: Erlbaum.

Ruscher, J. B., & Hammer, E. D. (1994). Revising disrupted impressions through conversation. *Journal of Personality and Social Psychology, 66*, 530-541.

Smith, V. L., & Ellsworth, P. C. (1987). The social psychology of eyewitness accuracy: Misleading questions and communicator expertise. *Journal of Applied Social Psychology, 72*, 294-300.

Stasser, G., Taylor, L. A., & Hanna, C. (1989). Information sampling in structured and unstructured discussions of three- and six-person groups. *Journal of Personality and Social Psychology, 57*, 67-78.

Vollrath, D. A., Sheppard, B. H., Hinsz, V. B., & Davis, J. H. (1989). Memory performance by decision-making groups and individuals. *Organizational Behavior and Human Decision Processes, 43*, 289-300.

Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185-208). New York: Springer-Verlag.

Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology, 61*, 923-929.

Wegner, D. M., Giuliano, T., & Hertel, P. T. (1985). Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253-276). New York: Springer-Verlag.

Wicklund, R. A. (1989). The appropriation of ideas. In P. B. Paulus (Ed.), *The psychology of group influence* (2nd ed.). Hillsdale, NJ: Erlbaum.

Yuker, H. E. (1955). Group atmosphere and memory. *Journal of Abnormal and Social Psychology, 51*, 17-23.