Chapter 8

Collective Awareness and Control

The discussion of collective strategy focused on ways to analyze strategic behavior in groups. Thus, its perspective was external, that of a researcher — or someone else interested for a different reason — observing and characterizing patterns in collaborative activities. In this chapter, the focus shifts from outside to inside.

For an individual to perform a task by following a known strategy involves a substantial degree of self-awareness and self-control. The person must address not only the substantive aspects of the task but also his or her own thinking. For example, the individual must recognize goals, select among options, and test work produced against intentions. These actions and decisions generate a continuous sequence of questions. *How do I get this idea across? Does the sentence sound right? Is it consistent with what I said earlier? Which step in the process should I go to next?* If these questions rise to consciousness, they become part of the thought process that constitutes awareness. But they may not. The individual may shift mental activity in a reflexive, seemingly automatic way as if an underlying question had been raised, but remain unaware of any such question. In either case, the individual experiences a continuum of thought that is both the matrix for strategic decisions and, at the same time, is at least partially comprised of those same decisions.

Can an analogous continuum of thought exist within a collaborative group? If such a continuum exists within individual minds, then it obviously exists within the minds of the group's individual members. But I have a hard time conceiving of a continuum of thought for the group as a whole.

This poses a dilemma. Earlier, I identified as the ideal artifacts that are coherent, consistent, as well as clean and simple in their design. When an objective standard exists, the artifact should also be correct, relative to that standard. For convenience, let me refer to these characteristics when applied to the artifact as *intellectual*

integrity. An artifact that has intellectual integrity suggests that the mind that produced it was aware of the whole and successfully related all of the parts to some unifying concept. However, for many collaborative projects, the size and complexity of the task precludes total awareness by any one member. How, then, can groups produce the continuity of thought needed to give their work integrity?

For software development, several approaches have been developed to help teams produce systems that are coherent and internally consistent. One approach is to require groups to go through a prescribed sequence of stages, each tied to a tangible milestone. This may help, but it does not guarantee success. Design documents must be verified against code, and the two kept consistent with one another though the lifetime of the system. In practice, this has been hard to achieve, particularly for large projects of long duration. Because we do not yet have a viable model of the collaborative process and because we have few comprehensive studies of collaborative projects, we have no *principled* basis for believing that one method is better than another or under what conditions a given method will produce reliable results.

One approach that has tried to address the one mind problem is the *chief programmer* method. It views the efforts of a system development team as primarily supporting, and thereby increasing the productivity of, a single highly skilled programmer (Mills, 1968; Baker, 1972), analogous to the team that supports a chief surgeon in the operating room. A chief programmer is responsible for the architecture of the system a team is building, and he or she writes the critical parts of the code. Supporting programmers write individual modules designed by the chief programmer and code the less critical parts of the system. The team also includes an assistant chief programmer and other supporting personnel, such as a librarian and secretary. Because one individual is responsible for the architecture of the system and for overseeing the work of subordinates, the design of the system is thereby held in a single mind — the chief programmer's. The problem with this method is that it does not scale to teams larger than six or eight members or to large projects that require hundreds of programmers and, hence, multiple teams.

Thus, neither ignoring the one mind problem or trying to solve it in a literal manner have proved satisfactory. A different approach is to formulate a concept of *collective awareness* as an abstraction and then try to use that abstraction to find ways of enabling groups to

function more coherently and more consistently. One way to do this is to look at the problem from an operational point of view.

Conceptual processing in individual human minds takes place under executive control. A metacognitive process appears to run in human short-term memory, making decisions about intentions and, in turn, activating specific cognitive processes — such as memory accesses, relational operations, and decision-making functions — that carry out those intentions. When a particular functional process completes its task, control returns to the executive process. Of course, the executive process may never be entirely passive, because it can interrupt functional processing when some unusual circumstance occurs that requires immediate attention, like answering the telephone or responding to a fire alarm. The goal, then, is to identify a comparable executive process that operates — or could operate — in groups.[7]

In the remainder of this chapter, I examine issues of awareness and control from a functional point of view. More specifically, I differentiate among several kinds of awareness that exist within collaborative groups in order to identify, albeit in a limited and abstract form, a concept of collective awareness that may help groups achieve a degree of integrity in their work comparable to that sometimes achieved by a single good mind working alone. I then look briefly at issues of collective control, analogous to individual self-control.

---

7 Metacognitive issues, such as awareness and control, have not received extensive attention from the cognitive science community. One study of note is Klatzky (1984). With respect to distributed cognition, see Rumelhart and McClelland (1986) for a discussion of fine-grained operations, but not large-grain problem-solving and knowledge-construction tasks, as considered here.

Awareness

Awareness in the sense that is considered here is an analog of self-awareness. Human beings are aware of past experiences — their own but also those of others they learn about and share vicariously — by virtue of their access to their respective long-term memories. Individuals are also aware of their own existence and their own thinking. Two analogous forms of awareness can be identified for collaborative groups: awareness of the group's long-term memory and awareness among the members of one another. Let's begin with the first.

The collective long-term memory has two parts: the artifact and the body of shared intangible knowledge. In chapter 6, the two were treated separately; with respect to awareness, they should be considered together. Although many of us would admit the possibility that intelligent agents may eventually traverse the artifact with some degree of awareness of its structure and content, for now any such artificial awareness is too limited to be relevant for a concept of collective intelligence. Consequently, I assume that awareness of the artifact will exist only in the minds of the human beings who comprise the group. Thus, awareness of the artifact is part of the group's intangible knowledge.

When we consider the level of awareness required for a concept of collective awareness, we should not set a goal that is higher than necessary or is higher than that found in individual human beings. None of us has total awareness of our respective long-term memories. At any one moment, we are aware of only a very small part of it — the contents of working memory. Although we may "sense" that we have accessed the very ideas we want for some purpose, we can never be sure that other, still more useful ideas are not stored in our long-term memories but have not been recalled. Thus, we can never be sure that what we are aware of is the most relevant knowledge potentially available to us. Rather, we activate and attend to portions of long-term memory and then move on to other portions. Thus, awareness of a large conceptual structure is a collage of partial awarenesses, generated at different times, at different granularities, and at different levels of abstraction. Consequently, we should not

expect collective awareness to be a total and complete awareness of the artifact, at least not an awareness that is activated at any one time.

Awareness in groups exists at several levels of detail. The most general is the body of intangible knowledge that is shared by all members of the group. It includes the overall goals of the project, its ways of operating, the strategies it uses to develop the artifact, its current status and problems, the relation of the project to the external environment, and so on. It also includes that comprehensive overview of the artifact held in common by the group. This awareness is not deep, if the project is large, but it provides each member with a sense of the whole.

At the other extreme is the deep, detailed, often technical, knowledge held by individual members. Depending on the project, a single individual is often responsible for a particular part. That person is expected to have deep, direct knowledge of the corresponding portions of the artifact as well as the underlying issues that inform them. Thus, the person's intangible knowledge goes well beyond what is present in the artifact to include reasons why the construct was built as it was, alternative designs considered and rejected, as well as more general knowledge of the task domain. The person uses this depth of knowledge to produce new components that become part of the artifact. Thus, the level of awareness and expertise required to generate a segment is significantly greater than that required for another person to understand it.

Between these extremes of general, shared knowledge and deep, individual generative knowledge is an intermediate level. It represents a form of awareness that is not often recognized in collaborative groups but may, ultimately, be the most important with respect to the integrity of the group's work. It is *thick* knowledge of adjacent or nearby areas. It takes the form of understanding, rather than generation. Thus, it is shared with the individuals or the team responsible for developing other parts of the artifact, but it is not as deep as their knowledge nor is it shared with the entire project. The most important function of thick knowledge is to inform work in the person's or team's own area by providing a kind of peripheral vision that extends into nearby areas. Thus, it provides a context for the interfaces between areas.

Developing this type of ancillary knowledge is important because no specification — formal or informal — can anticipate all of the issues that will arise in detailed design and implementation. Indeed, a

specification must be ambiguous relative to its implementation, otherwise it could be executed directly and its abstract expression would include as much detail as its implementation. Consequently, designers and programmers must interpret specifications, choose among alternative algorithms, make assumptions about processes that operate on the other side of interfaces, and otherwise exercise judgment. Bill Curtis and his colleagues at MCC (Microelectronics and Computer Technology Corporation) discovered that one of the scarcest and most important resources found (or not found) in industrial software projects across a number of different organizations is knowledge of the application area, which informs global design decisions (Curtis, Krasner, & Iscoe, 1988). Such knowledge can be considered a form of thick knowledge that extends across the interface between project and potential users and provides a context for requirements and specifications.

Thick shared knowledge can be developed through informal interactions, such as conversations, but it can also be developed through more formal mechanisms, such as reviews. Let's focus on reviews, because they can be institutionalized more easily than informal contacts. A technique used by many software projects as well as other kinds of projects is a type of formal review called a *structured walkthrough* (Yourdon, 1989). When a segment of a design document or a segment of computer code is available, the individual or team responsible assembles a group of colleagues, distributes copies of the relevant sections, and then during a structured meeting "walks" the group through the document or code. The primary goal is to animate the ideas or content for the portion of the artifact under review in the minds of the participants so that they can point out problems not foreseen by the developers.

Although a walkthrough is a time-consuming process, it requires far less time to review a portion of the artifact than it does to create it. Structured techniques make it possible to review a 20–30 page document or a comparable module of code in an hour or two of meeting time plus a similar amount of preparation time. On the other hand, generating the material to be reviewed is likely to take several orders of magnitude more time. Thus, we can identify a *slow build/fast review* cycle for creating and for comprehending portions of the collective long-term memory — the artifact — that is analogous to the *slow write/fast access* cycle identified by Newell to differentiate between the time required for encoding and storing concepts in human long-term memory and the time required for accessing them.

If we step back and look at the project as a whole, we can see that expanding intangible knowledge through slow build/fast review cycles can help members of a group develop fields of awareness in which their own knowledge is deep and concentrated over a relatively small portion of the artifact but extends outward with less, but still substantial, depth over a much larger area. Their respective fields of knowledge overlap with one another at the edges. These borders of *thick shared knowledge* can help members in adjacent teams keep their respective parts of the artifact consistent with one another. If all the teams or individuals in a project develop this type of expanded peripheral awareness, the group as a whole will have built a segmented, but overlapping awareness of the entire artifact.

This discussion of awareness has necessarily been abstract. Let me end it with a metaphor. My reading lamp is a hanging lamp with a handmade paper shade that resembles a bell jar (a cylinder with a rounded top). It is 14 inches both across the diameter of the cylinder and in height. Irregular pieces of paper, each 8–10 inches across and in muted shades of cream, pink, and tan are glued together to form the shade. The pieces do not abut one another; rather, each overlaps a half-inch or so with its neighbors. The overall shape is smooth and regular, enclosing a recessed bulb that casts light through the open bottom of the cylinder, and the whole thing glows softly.

Think of the individual awarenesses of group members as analogous to the individual pieces of paper in the lampshade. The boundaries of individual knowledge do not abut; rather, they overlap with one another. Thus, they share a common border — a boundary that is not a line, but has significant area. When this is the case, a portion of the artifact developed by one member will be informed by that individual's thick shared knowledge of nearby segments being developed by colleagues, and vice versa. In addition, portions of the artifact larger than a given individual's primary responsibility will have been held in that person's mind. Thus, the one mind condition will apply to a set of overlapping segments of the artifact.

When all the partial thick awarenesses are assembled, they form a whole that encloses the artifact. It is important that the surface of this collage be continuous, so that all parts of the artifact are surrounded. When this is the case, the artifact will have been subjected to the one mind condition, albeit in a piecemeal but continuous fashion.

Thus, we can identify three forms of awareness with respect to a group's long-term memory: close, detailed awareness of particular

segments of the artifact; less detailed, but still substantial, awareness of adjacent parts of the artifact; and the much thinner awareness of the artifact as a whole that is shared by the entire group. A different kind of awareness is the awareness members have of one another.

In addition to awareness of our respective long-term memories, we are also aware of our own thought processes. We can recall earlier instances when we thought about a particular problem or issue. At least occasionally, we may observe our thinking as it occurs to see that it is momentarily distorted by emotions or other socially induced factors. We also have a general sense of what we know and do not know, areas in which our knowledge is deep, and those in which it is not. These are different forms of self-awareness in which we briefly seem to step out of ourselves yet observe ourselves as a functioning mental process. Comparable forms of awareness exist within groups. Within the terms of this discussion, we can consider the problem with respect to the collective processor and the awareness one processor has of the other processors and of the system as a whole.

One of the primary reasons for assembling a group is to assemble the expertise required to carry out a project. For complex tasks, not all of the required expertise will be found in one head. Although it is conceivable that the work of the group could be partitioned so that requisite expertise is always matched with assigned task, this is seldom the case. Often an individual must call on his or her colleagues for help. The issue, then, is providing the group as a whole with a collective awareness of its members' respective specialized knowledge and expertise. Some groups refer to an individual with specialized knowledge as a *guru* in that area. Thus, an extremely valuable resource for a group is shared knowledge of who is a guru on what. Some computer systems, such as UNIX, even have built-in facilities for recording and accessing information on gurus. But, regardless of the mechanism, a group's knowledge of its gurus is comparable to an individual's knowledge of his or her own depth of knowledge in specific areas.

Another form of awareness is the awareness at any given moment one processor has of the other processors. For example, one member of the group may be aware, or wish to know, that another member is working in a nearby part of the artifact. This behavior is monitored at a very low level by the collaboration support system in its concurrency control mechanisms to insure that two members do not

try to change the very same part of the artifact at the same time. However, these mechanisms do not prevent one member's access from blocking that of another or of one member's subsequent work affecting earlier work done by another. Higher level interactions such as these must be permitted, but groups may also need help in monitoring domains of activity. For example, members may want tools that can provide a visual image of the artifact and show them where they are working within it. They may also want to see where colleagues are working. They may even wish to see a display over time of the "tracks" left by colleagues.

A third form of awareness involves the interaction between social and intellectual processes operating within the group. I discussed the fine-grained part of this issue in chapter 6 with respect to group-mediated cognition cycles. There, social and conceptual actions were seen to interleave as groups construct and/or use shared knowledge. A set of larger-grained issues has to do with the more overt effects members of the group have on one another and on their collective work. It would be nice if groups were purely intellectual organisms. But they are not. Tensions exist; friction occurs. These developments are inevitable. For the most part, they remain at the level of distraction, but they can become more intense and affect conceptual work. For example, one member may oppose an idea voiced by another not because the idea is bad but because of who said it. The opposite condition — supporting an idea because of friendship or attraction — is equally bad. These so very human situations are unlikely to go away, but a group should be aware of them and through its control and decision making procedures, try to insure that the integrity of its work is not compromised by them.

Thus, a group should be aware of itself as a dynamic, functioning organism as well as be aware of the artifact it is developing. From an experiential point of view, I still cannot envision a single, integrated awareness for a collective intelligence, but I can imagine it as a structural entity in the form of a collage of partial, but overlapping partial awarenesses. I can imagine how this form of collective awareness might function. And I can imagine how we might develop methods and tools to support it and to help groups develop it.

At the end of chapter 4, I briefly discussed an objection raised by Newell to the concept of collective intelligence, based on the limited rate at which human beings can transfer knowledge from one to

another. Newell asserted that for a group to behave as a coherent rational agent, each member of the group must know everything all of the other members of the group know. He is no doubt right that this is an impossible condition, but it may be too strong a requirement. Although not the complete knowledge Newell calls for, this composite of deep, comprehensive, connected awarenesses that comprise what I have called collective awareness may be sufficient to produce artifacts that are coherent and internally consistent to a degree comparable to those produced by a single good mind. If this is the case, then we can say that for practical purposes, the group has achieved a form of collective intelligence and thereby met Newell's objection.

Control

Control within groups has been studied from a variety of perspectives, including organizational theory, interpersonal relations, the characteristics of effective leaders, the impact of technology, patterns of communication, and the dynamics of groups over time. Much of that research focuses on social, as opposed to intellectual, factors. In chapter 1, I constrained this discussion to intellectual tasks. Thus, although the social dimensions of collaboration are important, for the kinds of groups being considered here, intellectual behavior is fundamental. Consequently, in this chapter I will consider control from a perspective in which conceptual and social dimensions are merged.

The form of control that is discussed is an extension of *self-control*. Self-control is an executive, metacognitive function that monitors the behavior of an individual — both mental and physical — and adjusts it in accord with some structure of goals, self-image, and/or set of external conditions. Thus, it monitors and responds to mental processes that lie below consciousness, to outside stimuli, and to properties and processes in the physical body. Although it is closely related to rational processing, it is not, itself, entirely rational.

I refer to the analogous form of control within a group as *collective control*, consistent with the terminology used in earlier discussions for other parts of a collective intelligence; however, I do not mean to imply that this form of control is necessarily democratic.

Collective control includes two large components: an organizational component and an intellectual component.

The organizational component is concerned with the group's operation and the procedures it uses. It is concerned with establishing an overall strategy for the group, setting priorities and goals, monitoring progress, and resolving organizational conflicts within the group. It is also responsible for obtaining the resources needed by the group and for interacting with the outside world. While essential, the organizational control function does not directly engage the substance of the group's work or directly manipulate the artifact. Consequently, it is similar to the executive function that monitors and controls what Vygotsky referred to as "lower" mental and physical processes.

The intellectual component does directly engage conceptual substance, because it is concerned with building a coherent, consistent structure of ideas. Consequently, this form of control places a high priority on "getting it right," recognizing validity as a primary requirement for the group's work. It may also try to achieve intellectual elegance — not as an end in itself, although clean, simple conceptual structures are often compelling, but because work with these characteristics is easier to understand, to communicate, and to maintain. This form of control includes "higher" mental functions, as Vygotsky used the term, including establishing a basic set of terms and concepts, constructing an overarching conceptual framework, and expanding and implementing that construct. Throughout this process, the control process monitors conceptual construction in order to modify the overall artifact design and to reconcile inconsistencies and differences of opinion that arise. As a result, it is responsible for the evolution and integrity of the artifact.

Collective control is, thus, the union of two types of executive functions — one organizational, the other intellectual. It is an abstraction that becomes actual in both the formal and the de facto control structures that exist within groups and in the individuals who function as leaders and/or occupy positions of authority. Thus, collective control can be achieved by a number of different organizational structures and styles of leadership. My own experience suggests that although a strong leader invested with authority is required, most often he or she leads best who leads least. If we tap on this simple maxim, it unfolds into the much larger concept of collective control I am describing.

First, intellectual integrity is best achieved if all members of a group *try* to achieve it in their individual work and in the areas where they have awareness and responsibility. In this way, they care what happens. Thus, developing a sense of shared ownership and responsibility throughout the group is important.

Second, all members must feel they can speak freely on substantive issues. Members doing detailed, technical work are often in the best positions to monitor and report problems. Those with leadership responsibilities, at all levels, must listen and respond. Otherwise, information will not flow freely, and the work of the group will be "brittle." The Challenger and Three Mile Island disasters were both dramatic failures caused by brittleness, in the sense that crucial information either did not flow or was not attended to across boundaries in the system. Less dramatic failures occur all the time when groups fail to achieve collective self-control.

Third, those who function as leaders perform acts of selection as often, if not more so, as they perform acts of generation. People in positions of authority do not have a franchise on good ideas; in fact, just as those directly responsible for detailed substantive work are in the best position to see problems, they are also the ones most likely to see new possibilities. However, those in leadership roles are often in a better position to select among new ideas, including testing them against the overall structure of the artifact and working out inconsistencies caused by adopting them. As the conceptual structure evolves, they are also responsible for articulating a new view of the whole to update shared intangible knowledge in the group.

Thus, the role I have sketched for a group's leaders is more a matter of perspective than authority, although authority must ultimately be vested in those leaders. In an organizational structure that has line authority, a team leader who works with several individuals is in a position to look over all of their shoulders and see how an idea generated in one context affects other contexts. It is conceivable that this same function of selection and reconciliation could be done by the team, itself, operating as a committee of the whole and without a designated leader. They could use their collective experience and knowledge to evaluate new ideas generated by any one of them and collectively make decisions that affect their part of the artifact.

Where this form of control breaks down is when there is disagreement in the group. In those instances, making a decision by

vote in a project that is primarily intellectual is not an acceptable solution; the majority may simply be wrong! The reason for this is the fact that the artifact is a material object that the group is responsible for building in such a way that it is coherent, consistent, and correct. Decisions that invalidate the integrity of the artifact are wrong, regardless of how they were arrived at and regardless of how many people agree with them. When consensus does not exist, some *one* individual must decide which option is (most) consistent with maintaining the intellectual integrity of the whole. Such a decision must be made on substantive grounds, not on the basis of its effects on social or organizational concerns.

Thus, the work of the group *may* happen by consensus; it probably *will* happen by consensus most of the time; but when consensus does not exist, some individual must have the authority to step in and make the decision that tries to maintain the intellectual integrity of the group's work, within the limits of that individual's capabilities. Thus, although I can admit ad hoc and network-based organizational structures that function most of the time across most of the group's activities, I cannot envision a group structure that can *reliably* produce work that is coherent, consistent, and, possibly, elegant that is not, ultimately, hierarchical, in the sense described here.

Finally, let me point out that collective control resembles, but is not identical to, the chief programmer model. It tries to achieve integrity and, perhaps, elegance in collaborative work by having the entire design or conceptual structure come together in a single mind — that of the overall project leader — but at a (possibly high) abstract level. However, it differs from the chief programmer model in two important respects. First, it assigns the leader a role that is much more integrative, based on selection and comprehension rather than on generation. By contrast, the chief programmer is expected to generate the primary architecture and important portions of the code. Second, it assumes that there may be other analogous leaders that serve similar roles with regard to individual teams. Consequently, the model can be scaled by including intermediate levels, consisting of groups of teams. If the project is large and includes multiple groups and levels, thick shared knowledge should overlap vertically between levels — just like it overlaps horizontally between teams on the same level — in order to provide continuity over the entire project. Thus, collective control, and the forms of project organization and integrated behavior it implies, achieves many of the benefits of the chief programmer model, without several important limitations.

In this chapter, I have looked at analogs for two of the more complex and elusive metacognitive functions that permeate human intelligence. The first is awareness. When we identify the characteristics of human intellectual work that we value most highly — coherence, consistency, correctness, and, elegance — it is difficult to imagine how work with these attributes could be produced without that structure of ideas having been held in its entirety by a single mind, if not actually produced by that mind. However, by considering awareness from a functional point of view, we may be able to construct mechanisms that can enable groups to achieve comparable results. One is a collage of partial but overlapping awarenesses based on thick, shared knowledge distributed over the group. Another is awareness within the group of the varied expertise held by its members.

The second issue considered was control. A mind capable of producing large artifacts that have intellectual integrity must also be disciplined. Not necessarily in a rigid way, but with enough self-control, informed by self-awareness, that it can test the constructs it produces against one another, against more general principles, and, perhaps, against some deeper aesthetics it has come to associate with "getting it right." The analog for self-control in a collaborative group is a collective control that balances hierarchical authority, required to resolve conflicts, with mechanisms that distribute responsibility throughout the group and generate vested concern for the integrity of their collective work.

Although it remains difficult to imagine how groups can achieve the same coherence and the same grace in their work that is sometimes achieved by individual minds working alone, we do not always have that option. Consequently, we must try to formulate mechanisms, such as those discussed here, that approximate the same functional characteristics within groups.

Issues for Research

A number of research issues emerge from concern for collective awareness and control. In the long term, we will need new tools and

new methods for studying collaborative behavior, such as those described in chapter 7, to make these problems tractable. However, it is not too early to begin addressing questions such as the following, because even partial results would yield significant benefits.

- *How do groups develop awareness of expertise distributed within the group and, in turn, use it effectively?*

Because most collaborative projects include individuals with complementary knowledge and skills, knowledge about who knows what must be distributed throughout the group. Also important is helping team members make good decisions about when to dig for knowledge on their own — from books, databases, etc. — and when to seek help from colleagues. A related problem is communication across knowledge boundaries. An individual seeking help may have difficulty finding a person who has the information he or she needs if the person does not know the terms in which to express that need so that a person who has the knowledge will recognize it as related to his or her expertise. Research in AI, library science, and automatic translation could be applied to this problem.

- *What are the properties and uses of thick shared knowledge?*

How is thick knowledge, developed by individuals with respect to adjacent areas of a project, built? How is it used? How much is needed? What happens if it is too thick or too thin? Can its development be institutionalized or is it a matter of individual choice and behavior? In what ways does thick "vertical" knowledge differ from thick "horizontal" knowledge?

- *How can we identify, analyze, and represent the collage of knowledge that surrounds the artifact?*

If a large project is surrounded, first, by a thin membrane of knowledge shared by the group as a whole and, second, by thick patches of more specialized knowledge, how can we characterize these bodies of knowledge? What specific concepts and structures do they consist of? How do the pieces fit together? How much should they overlap? Is it important for groups to share certain kinds of knowledge but not others? How can we tell if the artifact is completely surrounded? Does it matter?

- *What is the relationship between organizational structure — de facto as well as defined — and intellectual integrity in the artifact?*

If we characterize the relationships and structures that comprise collective control, can we then see how it affects the integrity of the group's work? Do properties in the artifacts correlate with specific patterns of behavior? If so, can we trace those patterns to particular control structures that generate or influence them? Do some organizational structures and procedures result in more cohesive work than others?

- *Which traits or experiences enable a leader to work effectively with structures of abstract symbols?*

Why can some people work with highly abstract symbols better than others? For large conceptual constructs, the architecture is formed at a level of high abstraction. Thus, each symbol or element in the design stands for a much larger component that, itself, may be quite deep. Consequently, high level symbols have long "tendrils" attached to them. If the design is to work well, those tendrils must descend gracefully and they must not get tangled with one another. Some individuals seem to have an intuitive feel for what lies beneath the symbols they work with, even when they have little direct knowledge of details at lower levels. They use that feel to produce designs that make clean separations that work well all the way down, or they use it to point out problems of decomposition in the designs of others. Is the skilled designer or project leader who works at a high level of abstraction necessarily someone who has worked his or her way up through the ranks and thereby developed this feel? Or, is it a native characteristic? Can it be learned? Can it be taught, either through training or mentor relationships?

- *What makes a conceptual structure elegant?*

What, exactly, does it mean to say that an intellectual product is simple and elegant? Is it only something we recognize when we see it? Or, can we develop parameters that will give us a more analytic sense of the characteristics that underlie such products? Is clean design related in some inherent way with the content domain? That is, are "seams" fundamental to a domain and, hence, discovered by a designer, or are they ultimately arbitrary and, hence, constructed?

How far down should an elegant design be expected to extend before it dissolves into arbitrary and/or messy detail?

- *What enables a group to produce elegant work?*

If it's easy for groups to design camels, why can't they learn to design eagles? Assuming we can arrive at a more basic understanding of what constitutes elegant design, how can groups achieve it? Is it a function of the individuals that comprise the group, its leader(s), the tools it works with, its procedures and organizational structure, and/or the environment in which it is located? Can groups learn this skill? If so, can we develop curricula and instructional programs to help groups develop it? One way to pursue this would be through case studies. If we can identify groups that have produced elegant products in the past, perhaps we could retrace their steps to see what enabled them to do so. Best of all would be to follow work in progress that turned out this way. What signals might alert us to such a group?