

## **Chapter 2**

# **Collaboration as an Information Processing Activity**

In most intellectual collaborations, the work of the group is ultimately focused in the production of some concrete product that constitutes successful completion of the project. Although groups differ significantly in their size, duration, the complexity of the products they produce, and the ways they go about doing their work, they also show surprising similarities in many of their activities, such as developing a body of shared knowledge about the task, using that knowledge to develop a plan of action, and, in turn, using the plan to guide their work. In this chapter, I illustrate both the diversity among collaborative groups as well as their similarity with respect to the types of information they use and produce.

The chapter is divided into two sections. First, three different collaborations are described to illustrate the range of behaviors commonly found in collaborative groups and to ground the discussion in concrete detail. Second, a simple information flow model is discussed; it describes the flow of information from one type to another, as concepts expressed in one form are transformed into a different form.

### **Three Scenarios**

The scenarios that follow describe three collaborative projects that range in duration from several weeks to several years. The groups range in size from 4 members to as many as 20. And their work results in an equally diverse set of products. In the first, an ad hoc group representing three different employment categories in a

## 12      2. Collaboration as Information Processing

university academic department work out a departmental policy for assigning parking permits. In the second, a task group in a federal agency is charged with planning and preparing the testimony a representative of that agency will give before a Congressional committee. The third scenario describes a larger group working in industry to develop a new software system along with its supporting documents. Although each scenario originated with an actual group or situation, I have taken liberties in the discussion to illustrate particular points. Thus, they should not be read as accurate descriptions of actual groups.

### **Scenario 1**

The first scenario concerns an ad hoc group formed within an academic department in a large university. Their charge is to draft a proposed departmental parking policy. The issue has recently become sensitive because of the closing of a large, nearby lot to make room for a new building. The group is composed of four members: representatives of the secretarial and technical support staff, the tenured or tenure-track teaching faculty, and the research faculty holding fixed-term appointments supported by outside research contracts, plus the associate department head who serves as chairperson of the group. They have 3 weeks in which to solve the problem, draft the policy statement, and report back to the head.

The group met right away to begin work. It quickly became apparent that parking was an emotional as well as practical issue for some of the members. The representative of the support staff told the group that his constituents thought spaces should be assigned solely on the basis of seniority. Because many of the staff had worked in the department for more than 10 years, this would mean that no desirable places would be available for new faculty to be recruited over the next few years, a position that was said to be unacceptable to the department head. The representative of the regular faculty argued for using seniority within employment category as the basis, with the faculty, research faculty, and staff categories falling in that order of priority. This was essentially the status quo. The research faculty representative proposed a rather complex formula that computed a number for each person based on points assigned for job category, rank within category, seniority, and several other factors. Although

the group remained relatively civilized toward one another, the first meeting ended with no solution in sight.

Over the next week and a half, the group met several times. At times, the meetings seemed productive; at other times, tempers flared. But, gradually, the group settled down and began to analyze the problem. Between two meetings, one member analyzed the data regarding lot requested versus lot assigned based on the current assignment scheme and concluded that the issue really boiled down to a shortage of eight "reasonably attractive" spaces that were desired by someone, but not available. Attempting to be funny, another member suggested that the department should rent those people spaces in a nearby commercial lot. After a few more wisecracks, someone asked the question: "Why not?" The group all looked at one another, paused, and then began to list the reasons why this could not be done. It would cost too much money. It would still mark some people as "second class citizens." And so on. But for each such reason, someone suggested a possible way around the problem. By the end of the meeting, they had not reached consensus, but no one seemed completely sure that the commercial lot option was impossible.

For the next few days, small knots of people could be seen talking earnestly among themselves, and one had the impression that most of the department was talking about the problem. When the team met several days later, the associate head reported that she had run a spread sheet on the costs and found that the difference between what the university charged for a parking space and what the commercial lot charged was only \$15 a month. Because the total came to less than \$1,500 for the eight spots per year, she felt it was financially feasible. Her remark was interpreted by the group to mean that she had also discussed this option with the department head and that he was willing to find the necessary funds. The staff representative still objected to the status implications, but he was having increasing difficulty arguing this position. Sensing this, he finally agreed to go along with the commercial lot option if the department would upgrade the workstations of people voluntarily accepting a less attractive lot assignment than they were entitled to. Because the department planned to upgrade all workstations over the next 2 years, the associate head said she saw no reason volunteers could not be given priority. With this, the group sensed that they had an agreement that, if not entirely to everyone's satisfaction, was at least workable.

While two members chatted, the other two drafted a short three-paragraph report, which they all then briefly discussed. After making

## 14      2. Collaboration as Information Processing

several minor changes, the group asked the chairperson to produce a clean copy of the report and submit it on their behalf to the department head.

In this scenario, the product produced by the group consisted of a one-page document written at the very end of the process. However, the size of this document does not accurately reflect the work of the group. A lot of information flowed back and forth within the group during their discussions. It also flowed between the group and the other members of the department in the form of both informal conversations and more "official" contacts such as that between the group chairperson and the department head. Most of this information remained intangible — in the heads of the participants — but it was crucial to the success of the group. The matter of upgrading workstations, which was not really related to the parking issue, addressed a social, rather than a conceptual, aspect of the problem. But, it, too, was crucial in making the solution palatable to an important constituency of the department.

Thus, the behavior of the group must be understood in both social and conceptual terms. Social aspects included both behaviors within actual group meetings as well as several different forms of interaction between the group and the rest of the department. Conceptual aspects included both tangible and intangible forms of information. The amount of intangible information generated over the 3 weeks was large compared with the tangible product finally produced by the group — the policy statement that was their charge. But, clearly, both forms of information were needed, and, for this task, probably in this proportion.

### **Scenario 2**

An investigative agency of the federal government has been asked to testify before a Congressional committee on potential use of the military in domestic drug enforcement. The agency has standing working groups concerned with the military, drug-related issues, and legal issues — each in a different division of the agency — but it has no established group concerned with this particular combination of issues. As a result, a five-person team — comprised of three area specialists, the person who will testify as the representative of the

agency, and her chief assistant — is assembled to plan and develop the testimony. The group is responsible for writing the approximately 15-page statement that will be entered into the official record of the committee and any visual aids or supporting materials that will be used by the witness during her testimony. Just before the hearing is held, the witness will plan her own brief (typically, 5-minute) oral presentation, based on the written statement developed by the group. The witness is the team leader and has authority over the work of the group with respect to this project. The group has approximately 2 1/2 months in which to complete the task. Although the assignment is regarded as important and high-profile, team members all have other responsibilities and assignments within their respective divisions that will also require their attention throughout the project.

The group began by reviewing previous work done in each of the three areas involved as well as information available from outside the agency. This task was divided among the members, following the first organizational meeting of the group, and occupied them for the first 2 weeks of the project. After that, the group met as a whole two or three times a week over the next several weeks. During these discussions, key points in the source materials that had been gathered were identified that might be relevant to the testimony. The group also compared notes on what they knew about the members of Congress on the committee and their known concerns and idiosyncrasies. Group members took turns listing points on a whiteboard, drawing various conceptual diagrams, and so forth. They also took turns writing up notes of their discussions and distributing them to the others.

Fairly early, the group crafted a one-sentence “message” to serve as the focal point of the testimony. This turned out to be a surprisingly difficult task. They could all identify important issues, but seeing which larger point they all added up to proved to be harder than they had anticipated. The message was also revised several times over the course of the project.

By the beginning of the fifth week, the group felt it was ready to shift gears and begin planning the actual written statement. During a particularly long meeting, they hammered out an outline that included the message, five major sections for the testimony statement, and the main points to be made in each. At the end of the meeting, they divided the sections up among themselves; most were assigned to a single individual, but two sections were assigned to two-person teams.

Throughout these discussions — and, in fact, throughout the project — the individual area specialists on the team met informally with their respective deputy division directors to keep them informed regarding the direction the project was taking. Occasionally, they would receive instructions regarding positions or arguments on an issue that were regarded as sensitive by the divisions. Although members would occasionally report these conversations, more often these concerns were simply absorbed into the perspective of the particular team member and voiced during the group's discussions. In general, the divisions made relatively few attempts to influence the message and strategy the group crafted, although they expected to be kept informed.

Over the next week, each individual or team was responsible for fleshing out additional levels of detail for the plan by identifying the specific facts or policy positions to be discussed under each point. At the end of that time, they again met as a group. During this meeting, they put their various pieces together, reviewed the whole extended plan, and revised it; they also assigned the task of looking further at several particular issues to the individual or team involved. At the end of the meeting, writing assignments were decided. One person was also given the task of collecting the final revised pieces of the plan and distributing copies of the whole to the group. Thus, by the beginning of the sixth week, the group was ready to begin drafting in earnest.

As group members wrote their individual sections, they referred to the overall plan for context — for example, one person discovered that an issue that had been viewed as a matter of policy was actually a point of law. Because legal issues were described in an earlier section being written by another member, this writer decided to meet with the other person to negotiate a change in the plan. The other writer agreed, and they sent a note to the rest of the group telling them of the change in the document plan. Several similar changes were negotiated by other members over the course of the project.

When drafts of all of the sections were complete, some 2 weeks later, they were assembled into a single document, and copies of the whole were sent to all five team members. Each reviewed the document alone before they met as a group to discuss it. It was then 3 weeks before the scheduled date of the hearing, and the document was rough. A lot of material was repeated, and the writing styles of several members were noticeably distinct. After some discussion of the whole document, the group went through it section by section, discussing potential problems. After two rather long meetings, they

completed their detailed review and had marked a number of places for further work. Again, the individual members of the group worked alone to revise their sections in accord with the discussion and the marked-up document. As the sections were revised, they were again circulated, marked, and returned to their respective authors. A “final” group draft was ready 10 days before the hearing.

At this point, they met with someone from the agency’s Congressional liaison office to discuss the actual hearing event and to identify questions expected from the various members of the committee. After the liaison person left, they also discussed visual aids for the presentation. They went back and forth on the issue, but finally decided to prepare two large (30 in. x 60 in.) poster boards to be placed on easels at the side of the hearing room during the testimony. One succinctly stated the “message” and three supporting points; the other showed in graphic form the results of an econometric model that predicted different impacts on availability of illegal drugs that would result from different levels of interdiction by the military. The order for the poster boards was placed with the agency’s graphics department.

Following this meeting, the witness who was also the team leader did a complete editing pass through the testimony statement for consistency and to make the language closer to her own writing/speaking style. Two days later, that version was passed to the head of the agency for his comments and, they hoped, approval. He liked it and had only minor changes, which were made. The statement was ready a full 48 hours before the hearing!

This collaboration is quite different from the first one. Part of this difference can be accounted for in the size of the products developed. Because the document that represented the goal of the project was longer, the group produced a number of intermediate products that were never intended to become part of the final document but served important instrumental roles. These included the notes and diagrams on the whiteboard, notes of meetings, and the different versions of the plan for the document. The intended product — the testimony statement — also went through several versions. Thus, the Scenario 2 task was more of a knowledge-construction task than a problem-solving task, as was the case in Scenario 1.

Although a great deal of important work went on in group meetings, the majority of time was spent in individual work —

examining resource materials, building and extending portions of the plan, drafting sections of the testimony statement, reviewing and revising drafts, and so forth. Although the group had to consciously attend to issues of internal consistency within the statement, this was not a particularly difficult or time-consuming problem for a document of this size. On the other hand, defining the overall message, particularly one that represented the consensus of the group, proved more difficult than they had anticipated. Finally, this group was more intellectual and less social in its emphasis, although like Scenario 1, Scenario 2 team members kept their constituencies informed as to the approach the group was taking. But, overall, the work of this group was far less emotionally charged than was the case for the group described earlier.

### **Scenario 3**

The third scenario concerns a group working within a software development department in a large computer company. The group, which ranged from 15 to 20 people during the project, collaborated over a 2-year period to develop a system intended to be used within the corporation to support a new "quality management" initiative that was to begin soon. However, if the initiative proved successful, the company planned to market system, training, and consulting services to other organizations starting similar programs. Because this project is too large to describe in as much detail as the first two scenarios, I focus on ways in which it was similar and dissimilar to the other two.

For most of the project, the members of the group were divided into three technical teams responsible for developing the three main components of the system — the underlying database, the user interface and its function, and communications. These teams ranged in size from three to six members. The group also included a project leader, assistant project leader, and two clerical support staff. The project leader designated one senior person in each team as team leader; these same three staff members plus the assistant project leader also functioned as a technical advisory committee to the project leader.

Much of the group's work was concerned with writing several design documents that provided different perspectives on the system the group was developing; hence, the process they followed was often similar to that of Scenario 2, only repeated several times. Scenario 3



was different, however, in its greater reliance on diagrams and in producing an altogether different kind of "document" — the computer code that constituted the system, itself.

The (conventional) documents they wrote over the course of the project began with an initial "whitepaper" that provided a conceptual description of the system. It was written by the team leader before the project began. It also described the rationale behind the quality management initiative and listed a number of information management and access functions that would be needed if the initiative was to be successful.

This list served as the starting point for a second document — a more formal requirements statement that identified the specific computer operations the new system would support as well as other technical requirements for hardware, performance, the programming languages and toolkits to be used, and so on. The requirements document was written collaboratively by the project leader, the assistant leader, and representatives from several departments expected to use the new system.

A third document described the abstract design or architecture of the system. It was developed iteratively over a 3-month period. Overall design for the system was discussed intensely for the first few weeks in meetings attended by representatives from all three teams, then off and on for the next several months, and then, for the remainder of the project, only occasionally to resolve problems that arose. Once the general architecture became relatively firm — after the first month — the more detailed designs for the different parts of the system were developed by the individual teams. As the various segments were completed, they were assembled into a single document. Throughout this process, the project leader's technical advisory committee met to review the individual parts as they became available.

A more detailed specification of functions and interfaces supplemented the architecture document. Work on this document overlapped work on the architecture. These descriptions exhaustively defined the interfaces for each module as well as the functions included within them. As the architecture document neared "completion" and work on it tapered off, work on the specifications increased. However, both the architecture and the specifications remained incomplete throughout the project. As the system was coded and tested, the teams often found that changes were needed in the

system design. Consequently, they updated the architecture, specifications and, sometimes, the requirements documents throughout the project. A larger problem was created, however, when they did not make these changes. This problem arose primarily from programmers making design changes in the system code but not updating the design documents to reflect those changes. Thus, toward the end of the project, the specifications and the other design documents became less and less consistent with one another and with the computer code that comprised the actual system.

A separate document outlined a set of tests that were used during development to test individual components and, near the end, to test the system as a whole. One member from each team was given responsibility for writing the test plan for his or her team's part of the system. They, along with the assistant project leader, also functioned as a test group, particularly toward the end of the project.

The last major document was a set of user instructions. An initial version was actually written quite early in the project as part of the specification for the user interface, but it was barely more than an outline of user operations. Near the end, two members of the interface group were asked to rewrite this document to include more rationale and explanation. It took them approximately 2 months to produce a new draft; however, the final user documentation was not completed until well after the development project ended and then only after a technical writer rewrote the draft from scratch.

Thus, the project produced a number of "formal" documents that included the concept paper, requirements, architecture, specifications, test plan, and user documentation. They also produced a number of "informal" documents; these included notes of meetings, memos and letters, various project plans and time-lines, progress reports, a budget, and brief explanations included in the computer program itself. More ephemeral were jottings on whiteboards, personal notes, and agreements and understandings that were reached but never documented. In general, the project took pains to produce and keep track of their formal documents, less so for their informal ones, and not at all for ephemeral products. As a result, people would occasionally search frantically for some lost piece of paper, but, in general, the group got along relatively well with the materials they kept, aside from the problems of consistency noted previously.

Throughout the project, diagrams played a more substantial role in the group's work than in the other two scenarios. Drawings of

various sorts comprised a relatively large part of the planning documents, particularly the architectural description. They were also prevalent in work at the whiteboard during meetings. The more abstract their discussions, the more important they became, particularly for ironing out differences in understanding of technical points and for exploring different design options. In fact, key decisions often turned on these diagrams, some of which were preserved in meeting notes but many of which were retained only in the heads of those attending the meeting.

The third major type of product produced by this collaborative group was the computer source code. Written in a programming language, this “document” constituted the actual goal of the project — the computer system the group was building. Consequently, all of the other products described previously were secondary, although writing the code without them would have been impossible. Just as the group produced several different kinds of documents, so they produced several different kinds of computer code.

Prior to the beginning of the project, the project leader and the vice-president in whose division the project was situated had decided that the system would be developed using a relatively new approach to software design known as object-oriented programming. This approach was said by its proponents to result in systems that are more modular, easier to understand and maintain, and that future projects should be able to reuse portions of their code with relatively few changes. In picking the project personnel, the project leader was careful to select several senior programmers experienced in both object-oriented design and programming, and it was from this group that the three team leaders were named. However, most of the other programmers were skilled in more traditional methods. As a result, the project faced a substantial learning curve at the beginning to absorb the object-oriented paradigm and to get up to speed in the specific object-oriented programming language selected.

Training was supervised by the assistant project leader, but most actual instruction was done by the three team leaders. For the first several weeks, all three teams met as a group and were instructed by the team leaders, taking turns. After that, training merged with other activities and was done on a team-by-team basis. During both stages of training, the individual group members wrote a considerable amount of computer code that was never intended to be part of the system. At first, they wrote small bits of code in order to understand how a particular technical feature in the programming language

## 22      2. Collaboration as Information Processing

worked. After that, they wrote short programs to which they applied principles of object-oriented design. Later, they wrote code as a way of exploring and learning new toolkits and utilities. To some extent, exploratory programming continued throughout the project, but most of it occurred during the first 2 or 3 months and overlapped with early design.

The second type of computer code written by the project was the system itself. The project's general strategy was, first, to prepare a detailed design for a module as described previously and, then, to write the code. However, as development moved to more and more detailed levels of the architecture, designing and coding often merged. For example, early design did not always anticipate unusual exception conditions. When the teams became aware of issues such as these, they were forced to design and code at the same time, sometimes leading to inconsistencies between system design as documented and system design as implemented in the code. Once written, the code went through numerous "revisions" during debugging and testing. Gradually, the teams accumulated libraries of system modules that had been tested and were thought to be relatively free of errors. These were ultimately assembled to form the complete system. However, as problems were discovered and corrected, these libraries had to be updated accordingly and the system reassembled.

The third type of code produced by the project was code intended to assist with debugging and testing. Each team developed a set of test programs to simulate the functions of the modules their components interacted with, to provide an exhaustive set of test conditions, and to simulate different load conditions for evaluating performance. Although many of the test programs were written by the individual team members to test their own work, quite a few were eventually gathered into standard "test suites" used in the latter part of the project by the testing team to ensure that corrections made to fix one problem did not introduce new ones.

Several differences stand out in this scenario. Perhaps the most prominent one is the difference in scale. Scenario 3 was an order of magnitude longer in duration than Scenario 2, but it was more than two orders of magnitude larger in the total size of the products it developed. Whereas both Scenario 1 and Scenario 2 each went through one complete cycle of development for their respective documents, Scenario 3 went through many such cycles for the

different documents and modules of computer code developed during the project.

These differences in scale also led to differences in project organization. Most of Scenario 3's work was done in relatively independent subgroups — thus, the project was a group of groups. In contrast, Scenario 2 functioned largely as a single coherent group that periodically delegated tasks to individuals or two-person teams but quickly assembled their independent efforts into a whole. Thus, one could normally say that at a given time the group as a whole was engaged in a particular activity, such as brainstorming, planning, drafting, editing, and so forth. In contrast, the different Scenario 3 teams progressed at different rates through their respective work plans and merged their independent efforts far less frequently. Consequently, one could not say at any given time exactly which activity the group as a whole was engaged in, because they were often doing different things.

## Model of Information Type and Flow

The three projects described previously are quite different from one another in terms of scale, duration, task, and the intellectual activities they used to accomplish their respective goals. But if we focus on the types of information they used and produced, we can also see similarities.

Each of the three groups was concerned with producing some type of tangible product that represented successful completion of the group's primary goal. In Scenario 1, this target product was the brief written recommendation regarding parking policy they wrote during their final meeting. In Scenario 2, it was the 15-page written testimony statement that was included in the record of the Congressional hearing and the two supporting poster boards referred to by the witness in her oral statement. In Scenario 3, the target product was the computer system defined by the source code and the supporting documents, such as the user documentation, delivered with that system.

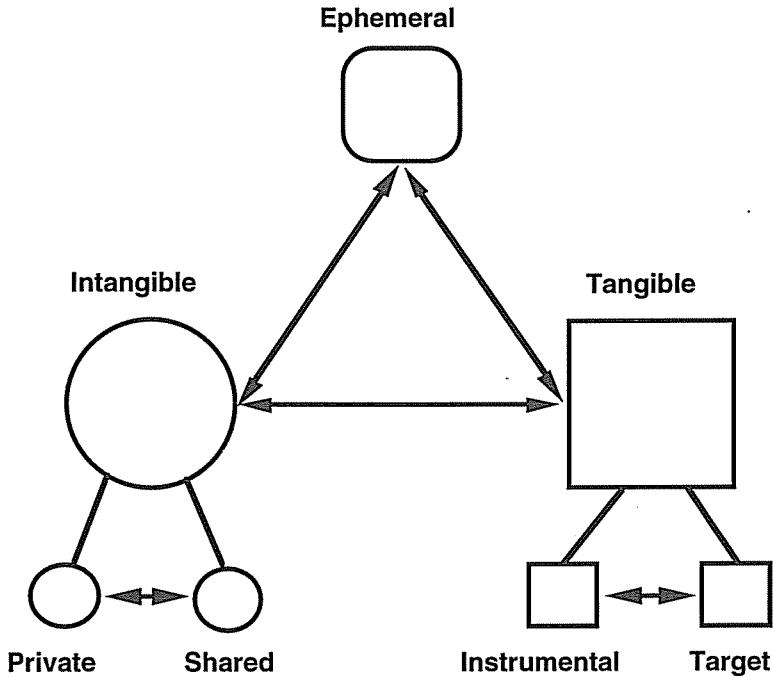
However, target products tell only part of the story. The groups also produced instrumental products to help them with their task.

These ranged from meeting notes and whiteboard jottings, to plans and outlines for the target product, to large independent documents, such as requirements and test plans. Some of these, such as the whiteboard drawings, were ephemeral. They came into existence for a brief time, served their purpose, but were then destroyed or lost. Others, such as the software design documents, were persistent; they were kept by the group and maintained with varying degrees of consistency from the time they were created until the end of the project.

In addition to these tangible forms of information, all three groups also used and/or produced several types of intangible knowledge. Intangible knowledge is carried in the heads of group members. In some cases, it was eventually incorporated into one or more tangible products. For example, in Scenario 1, the organizational change that was finally recommended was first discussed, along with several other possibilities, both inside and outside the group before it was finally encoded into written form. Thus, the information existed in intangible form before it was transformed into a tangible, written document. In other cases, information remained intangible throughout the project. This was true for the options discussed but not selected by the group in Scenario 1.

In some cases, intangible knowledge is shared by the group as a whole. For example, during early group discussion, individual members often voice their different understandings of the group's goal before the group eventually settles on a common approximation acceptable and known to all members. However, not all shared knowledge is developed through explicit group activities; for example, shared knowledge is also carried in the culture of the organization. In Scenario 2, all members of the group knew the general format of a Congressional hearing; what was not common knowledge were the idiosyncrasies of the individual Congress members on the committee.

In other cases, intangible knowledge remains limited to a specific individual or a subset of individuals and, thus, is private with respect to the group as a whole. This is frequently the case when members are selected because they have particular expertise or perspectives. During discussions, some of an individual's private knowledge may become shared as that person speaks to issues from his or her base of private knowledge, but much of this base remains out of sight from the group as a whole and comes into play only indirectly by informing that individual's work or views.




**Fig. 2.1.** Types of information produced by collaborative groups and the transformations from one type to another.

Thus, collaborative groups as a matter of course develop three basic types of information: *tangible*, *intangible*, and *ephemeral*. Tangible knowledge can be divided into *target products* that represent successful completion of the group's task and *instrumental products* that support the group's work on the target product but are not part of that product. The collection of target and instrumental products developed and maintained by a group during a project constitutes the group's *artifact*. Intangible knowledge does not take tangible form but, rather, remains within the heads of the members of the group. Some is *shared* to an approximation by all members of the group; other is *private* with respect to an individual or a subset of the group. *Ephemeral products* lie somewhere between tangible and intangible knowledge. This information is given physical form for brief periods of time, but unlike the instrumental products that are included within

the artifact, ephemeral products are destroyed or lost. These three basic types of information and their respective subtypes are shown in Fig. 2.1.

The collaborative process produces a *flow of information*, as information in one form is transformed into information in a different form. This flow is indicated in Fig. 2.1 by arrows between information types. For example, during group discussions, private knowledge held by one member may become shared knowledge held by the group if that person explains a privately held concept to the group. When this happens, information can be viewed as flowing from one individual to the group. If the individual with the private knowledge uses the whiteboard to draw a diagram to help explain the concept, the information is transformed from private to ephemeral to shared knowledge; Fig. 2.2 shows this particular flow. There are a number of different flows that can be observed in groups, produced by a variety of different processes and activities. It remains a task for research to uncover these.



**Fig. 2.2.** Flow of information from private to ephemeral to shared.

Collaborative groups produce a mix of products and knowledge, but different groups go about it in different ways, on different scales, in accord with different strategies. Consequently, groups can be characterized according to the nature of the artifact and the flow of information from one type to another they produce over the course of



a project. To see this, let's look back briefly at the three scenarios from the perspective of the information flow model.

In the first scenario, most of the group's time and effort was spent building a body of shared intangible knowledge, that is, analyzing the parking problem from the perspectives of the different constituencies and finding a solution acceptable to both the members of the group as well as those they represented. The actual target product consisted of only a short three-paragraph statement written at the very end of the project. Ephemeral products were also minimal and the group produced no instrumental products.

In the second scenario, the target product — the 15-page testimony statement and the two accompanying poster boards — was much more substantial, and direct work on it constituted a much larger proportion of the group's effort. This group also produced several different instrumental products — for example, meeting notes and the statement plan — that played important roles in the project. Scenario 2 undoubtedly developed a much larger body of shared knowledge than did the Scenario 1 group; however, this difference is unlikely to be proportional to the difference in size of their respective target products, which had an approximately 20:1 ratio. On the other hand, private knowledge played a much more important role in Scenario 2, because individuals and two-person teams, drawing largely on their individual areas of expertise, developed the detailed plans for the different sections of the statement and then wrote those sections.

The Scenario 2 group also went through a number of shifts between private and group work. This strategy permitted work to go on in parallel; but it also created problems of consistency, for example, when the group integrated the different parts of their detailed plan and, later, when they merged the draft sections of the testimony statement. However, because the scale of the project was relatively modest, resolving these inconsistencies was not particularly difficult or time-consuming.

The Scenario 3 project was several orders of magnitude larger and more complex than the Scenario 2 project. Its target product — the computer system — was a large "document" consisting of more than 1,000 pages of source code and accompanying comments. Instrumental products included a number of complete documents. The body of intangible knowledge shared by the group as a whole included knowledge of project goals, the large-scale architecture of the system, basic principles of object-oriented design, and a new programming

language most team members learned at the beginning of the project. Although the volume of shared knowledge for this project was larger than that of the other two, it was probably not proportionally larger relative to the difference in size of their respective artifacts. On the other hand, private intangible knowledge played a much larger role in Scenario 3 than in the other two scenarios, because individual teams possessed and developed specific technical expertise in their assigned portions of the system design. Finally, ephemeral products played a larger and more important role, particularly during the various design activities of the group.

Consistency requirements were more stringent and more apparent for the software development group. In Scenario 2, consistency relationships existed primarily within a single document, the testimony statement; however, in Scenario 3, consistency relationships also existed between documents, such as between the architecture document and the source code. Because of this and because different components were developed by different teams, keeping the whole artifact consistent was much harder. In fact, the group did not always do this, particularly toward the end of the project. These inconsistencies caused problems for the development group, and they can be expected to cause problems for those who will maintain the system in future years. Thus, as a project grows in size and complexity, maintaining consistency and coherence in the artifact is likely to become increasingly important and increasingly difficult, requiring more and more of a group's resources.

## Issues for Research

Implicit in this discussion are a number of unanswered questions and issues for further research. In this section, I briefly discuss several of these. This list, of course, is not exhaustive. Rather, it is intended to suggest the type of questions and issues that can be inferred from the discussion. I hope that others occurred to the reader, that they were noted, and that they will stimulate discussion and, perhaps, future research.

- *Develop detailed portraits of specific collaborative groups*

Detailed descriptions for a broad range of collaborative groups would be valuable for a number of reasons. First, they would open the area up for inquiry so that we could see what collaboration is as a general mode of work across different tasks. From this data we could also infer a repertoire of both common and unusual behaviors. And we could begin to see which behaviors and patterns of work tend to result in more effective or more efficient collaborations versus those that are less productive. Latour & Woolgar (1979) provided a good example of close description based on ethnographic methods; for computer-based collaborations, we should be able to supplement their methods with system recorded data to observe a greater range of individual and collective behaviors as well as the rhythms and interactions between the two.

- *Test and refine the information flow model*

The information flow model needs to be validated against a wide range of collaborations in different organizational contexts. As detailed portraits of groups become available, the model should be tested against those data and updated as needed.

- *Identify the vocabulary of transformation processes that occur in collaborative groups*

A related issue is further refinement of the model. Each of the transformations depicted as a line between information types signifies a wide range of more detailed processes. For example, the transformation from private to shared intangible knowledge frequently takes place in group meetings (it is not, of course, limited to such gatherings), but it occurs in different ways. It can come about through a statement made by an individual in a discussion, a statement within a more formal presentation, or a response to a question. Each such action consists of a complex structure of finer-grained social and intellectual processes. Thus, by a vocabulary of collaborative actions I mean a list of the different kinds of behaviors that transform information of one type into information of another type. Such a list will, of course, be extensive if it is to cover the many different situations in which collaboration takes place. But identifying them will be worth the effort because these actions could serve as the basic elements in a process model of collaboration.

- *Identify sequences of collaborative actions*

Collaborative actions normally take place in sequences or “phrases” rather than as isolated events. For example, a discussion normally includes a number of different statements by different members of the group, interspersed with individuals' drawing diagrams, referring to documents present or absent, showing transparencies, waving their arms in the air, and so on. These different events are woven into a fabric of discourse. If identifying a vocabulary of actions is a first level of analysis and characterization, identifying short repeated sequences of actions is a second level. Groups use such sequences to build small component structures and/or to link new information into an existing conceptual structure. Thus, if we are to understand how groups develop large, complex artifacts over long durations, we must first identify the basic process sequences they habitually use.