8

The Technology of Team Navigation

Edwin Hutchins
University of California, San Diego

Abstract

Large ships are navigated by a team of people working with a specialized set of tools. This chapter looks first at the nature of the tools and the navigation task and argues that the tools do not amplify the cognitive abilities of the team members, but instead transform what are normally difficult cognitive tasks into easy ones. It then considers the division of labor among the members of the navigation team and the techniques by which they coordinate their work activities. The progress of various team members through the career cycle of navigation practitioners produces an overlapping distribution of expertise that makes it possible for the team to achieve training and job performance in a single activity.

NAVIGATION TECHNOLOGY

Step aboard any modern ship and you will find yourself surrounded by an incredible mixture of technologies. Systems of levers, cables, and pulleys that Leonardo daVinci would have understood perfectly dominate the deck, while just inside a metal hatchway are sophisticated electronic systems whose design and operation are fully understood by only a handful of people in the world. Similarly, the technology of the navigation suite is an odd mixture of old and new. Every functional chart house contains a collection of artifacts in which one can see the long tradition of the practice of navigation at sea. In a typical hour of navigation work, a modern navigator is likely to use technological devices whose basic designs vary in age from a few years to several centuries.[1]

At all times while a naval vessel is underway, a plot of its past and projected movements is maintained. Day and night, whenever a ship is neither tied to a pier nor at anchor, navigation computations are performed. In a long passage, the navigation activities may be continuously performed for weeks or even months on end. Most of the time the work of navigation is conducted by one person working alone, but when a ship leaves or enters port, or operates in any other environment where maneuverability is restricted, the computational requirements of the task may exceed the capabilities of any individual and the navigation duties are then carried out by a team of individuals working together.

This chapter considers the technology involved in a navigation task called *piloting,* which is the principal method of navigation when ships are moving in restricted waters. Technologically, this is a particularly conservative practice. The gyrocompass is the most advanced piece of technology involved, and the principles of its operation were well worked out by the end of the First World War. The other technological devices, a telephone circuit for communication among members of the navigation team, a telescopic sight, ballpoint pens and log books, pencils and plotting tools, and the Mercator projection chart are hardly what anyone would consider examples of high technology.

Yet there is something about these tools that keeps them in use in this important endeavor in spite of the introduction of more advanced technologies around them. Without doubt, much of the piloting task could easily be automated with existing computer technology, and I am frequently asked, but find myself unable to answer, why the navigation team and its old tools have not yet been replaced by perhaps a single navigator using an electronic chart. The pursuit of an answer to that question would doubtless lead us to consider who wants such a change, how technology needs are recognized and promoted, who pays for the development of new technology, what other tasks have more urgent technological needs, what resources are available to proponents of competing schemes, and so on. Such an exercise would likely demonstrate once again that technological change is as much a social as a technical process. Instead of speculating on why this technology has not yet been overtaken by "something better," it seems more productive to look closely at how this existing technology is actually used.

I argue first that the old technologies that survive here constitute a different kind of advanced technology. The altitude of a technology might not be measured only in terms of the sophistication of the inner workings of the device hardware, but also in terms of the extent to which the device renders an important problem easy to solve. I then show that the use of these tools supports a distribution of knowledge among the members of the navigation team that makes the system very robust in the face of individual component failures.

Piloting is a specialized task that, in its ordinary operation, confronts a limited set of problems, each of which has a well-understood structure. The problem that confronts a navigator is usually not to figure out how to do the processing of the information in order to get an answer. That has already been worked out. Rather the problem, in most instances, is simply to use the existing tools and techniques to process the information that is gathered by the system and to produce an appropriate evaluation of the ship's situation or an appropriate recommendation about how the ship should proceed in order to get where it is supposed to go. This task therefore differs from many that are considered by other chapters in this volume inasmuch as it is a well-understood, narrowly bounded task that is performed in a technological environment that has been evolving for hundreds of years.

NAVIGATING LARGE SHIPS

Guiding a large ship into or out of a harbor is a difficult task. Ships are massive objects; their inertia makes them slow to respond to changes in propeller speed or rudder position. Because of this response lag, changes in direction or speed of ships must be anticipated and planned well in advance. Depending on the characteristics of the ship and its velocity, the actions that will bring it to a stop or turn it around, for example, may need to be taken tens of seconds, or even minutes, before the ship arrives at the desired turning or stopping point. Aboard naval vessels, a continuous plot of the position of the ship is maintained to support decisions concerning its motion.[2] An officer called the conning officer is nominally responsible for the decisions about the motion of the ship, but for the most part the conning officer does not make such decisions. Usually, such decisions are actually made by the navigation team and passed to the conning officer as recommendations, such as "Recommend coming right to zero one seven at this

---

[1] In what is perhaps the ultimate technological contrast, I once saw a junior officer take latitude and longitude coordinates from a satellite navigation computer, and roughly plot the ship's position on a chart using the span of his fingers as a measurement tool.

[2] Such complete records are not always kept aboard merchant vessels and are not absolutely essential to the task of navigating the ship in restricted waters. It is possible for an experienced pilot to eyeball the passage and make judgments concerning control of the ship without the support of the computations that are carried out on the chart. Aboard naval vessels, however, such records are always kept for reasons of safety primarily, but also for purposes of accountability so that if there should be a problem, the ship will be able to show exactly what it was doing and where it was at the time of the mishap.

time." The conning officer considers the recommendation in the light of the ship's overall situation, and if the recommendation is appropriate, will act on it by giving orders to the helmsman, who steers the ship, or to the leehelmsman, who controls the ship's engines.[3]

The navigation activity is event-driven in the sense that the navigation team must keep pace with the movements of the ship. Unlike many decision-making settings, when something goes wrong aboard ship, quitting the task or starting over from scratch are not available options. The work must go on. In fact, the conditions under which the task is most difficult are usually the conditions under which its correct and timely performance is most important.

In this chapter I consider the members of the navigation team together with the team's social organization and the tools of the navigation trade as a system of socially distributed computation. In particular, I consider the roles of people and technological devices in the computations that constitute the navigation task.

Position Fixing by Visual Bearings

In order to plan the motions of the ship, the navigation team must establish the position of the ship and compute its future positions. The most important piece of technology in this task is the *navigation chart.* A navigation chart is a specially constructed model of a real geographical space. The ship is somewhere in space and to determine or "fix" the position of the ship is to find the point on the appropriate chart that corresponds to the position of the ship in the world.

The simplest form of position fixing, and the one that concerns us here, is position fixing by visual bearings. For this one needs a chart of the region around the ship, and a way to measure the direction, conventionally with respect to north, of the line of sight connecting the ship and some landmark on the shore. The direction of a landmark from the ship is called the landmark's bearing. Imagine the line of sight in space between the ship and a known landmark. Although we know that one end of the line is at the landmark and we know the direction of the line, we can't just draw a line on the chart that corresponds to the line of sight between ship and landmark because we don't know where the other end of the line is. The other end of the line is where the ship is and that is what we are trying to discover.

Suppose we draw a line on the chart starting at the location of the symbol for the landmark on the chart and extend it past where we think the ship is, perhaps off the edge of the chart if we are really unsure. We still don't know just where the ship is, but we do know it must have been somewhere on that line when the bearing was observed. Such a line is called a line of position. If we have another line of position, constructed on the basis of the direction of the line of sight to another known landmark, then we know that the ship is also on that line. If the ship was on both of these lines at the same time, the only place it could have been is where the lines intersect. Each line of position thus provides a one-dimensional constraint on the position from which the landmark was observed. The intersection of two lines of position uniquely constrains the location from which the observations were made. In practice, a third line of position with respect to another landmark is constructed. The three lines of position form a triangle and the size of the triangle formed is an indication of the quality of the position fix. It is sometimes said that the anxiety of the navigator is proportional to the size of the fix triangle.

The Fix Cycle

The necessity for continuously plotting the ship's position, projecting the future track, and preparing to plot the next position is satisfied by a cycle of activity called the fix cycle. If the ship is near enough to land to take visual bearings on terrestrial landmarks, yet on a heading that will not soon bring it near shallow water or other dangers, the cycle may be completed at a leisurely pace of say once every half hour.[4] When the ship is in restricted waters, however, it may be necessary to complete the cycle on one-minute intervals. Under these conditions, no single person could make all of the observations and do all of the computations required to complete the cycle in the amount of time available.

The Fix Cycle in Sea and Anchor Detail. When the ship is operating in restricted waters, the work of the fix cycle is distributed across a team of six people.[5] The duty stations of the members of the team in the configuration called *sea and anchor detail* are shown in Fig. 8.1 as elliptical shapes. We can follow the fix cycle by following information through the system.

New information about the location of the ship comes from the bearing takers on the wings of the ship (positions 1 and 2 in Fig. 8.1). They find landmarks on the shore in the vicinity of the ship and measure the bearings of the landmarks (direction with respect to north) with a special telescopic sighting device called an alidade. The true north directional reference is provided by a gyrocompass repeater that is mounted under the alidade. A

---

[3] On older ships, of the sort aboard which this study was conducted, the leehelmsman operates a device called the engine order telegraph, which indicates the desired engine speed to personnel in the engineroom who actually control the engines directly by operating valves that admit high pressure steam to the propulsion turbines.

[4] Once the ship is out of sight of land, position fixes are made using a number of techniques beyond the scope of this chapter. These include radar, celestial, loran, omega, and satellite navigation.

[5] On other ships, and on this ship in different circumstances, the team may be somewhat larger or smaller depending on the availability of qualified personnel.

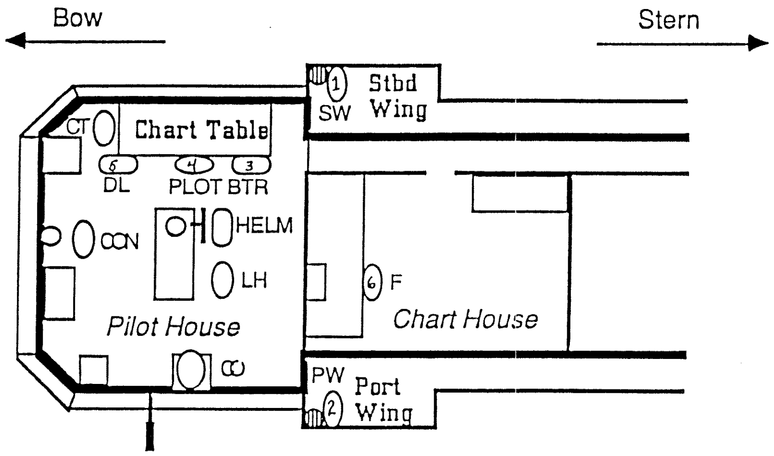

FIG. 8.1. Watchstander positions for sea and anchor detail.

prism in the alidade permits the image of the gyrocompass scale to be superimposed on the view of the landmark. (An illustration depicting the view through such a sight is shown in Fig. 8.2.) The bearing takers read the measured bearings and then report them over a telephone circuit to the bearing timer-recorder.

The bearing timer-recorder (position 3 in Fig. 8.1) stands at the chart table inside the pilothouse. This person talks to the bearing takers out on the wings and writes the reported bearings in a book called the *bearing log,* which is kept on the chart table.

The plotter, (position 4 in Fig. 8.1) plots the bearings that are reported by the bearing takers. This person normally has no direct communication with the bearing takers, but is either told the bearings by the bearing timer-recorder, or reads them out of the bearing log. From the perspective of the plotter, the bearing timer-recorder is an information buffer. In order to make a high-quality fix, the bearings should be observed as quickly and as nearly simultaneously as possible. It takes much longer to plot a line of position than it does to make the observation of the bearing, so the activities of the plotter and the bearing takers have different distributions in time. The activities of the bearing timer-recorder not only provide a permanent record of the observations made, they permit the bearing takers and the plotter to work, each at their most productive rates, without having to coordinate their activities in time. The bearing log is thus a technology that is used to "ease the constraints of time/activity match" (McGrath, chap. 2 in this volume).

In addition to plotting the ship's position, the plotter also projects where the ship will be at the time of the next few fix observations. This requires a knowledge of the heading and speed of the ship. The plotter normally reads these from the deck log, which lies on the chart table to his left.

The keeper of the deck log (position 5 in Fig. 8.1) maintains the deck log in which all events of consequence for the ship are recorded. All commands given by the conning officer to the helmsman (HELM in Fig. 8.1) concerning the course to steer, and all orders to the leehelmsman (LH in Fig. 8.1) concerning the speed to order from the engineroom are recorded.

When the projected position of the ship has been plotted, the bearing timer-recorder consults with the plotter to decide which landmarks will be appropriately situated for the next position fix, and assigns the chosen landmarks to the bearing takers by talking to them on the phone circuit. In choosing landmarks, the plotter and the bearing timer-recorder are looking for a set of three landmarks such that the lines from the three landmarks to the projected position of the ship intersect at reasonably steep angles. If the lines from any pair intersect at a shallow angle, then a small angular error in either one will move their point of intersection (one corner of the fix triangle) considerably, and add uncertainty to the fix. If the angles of the

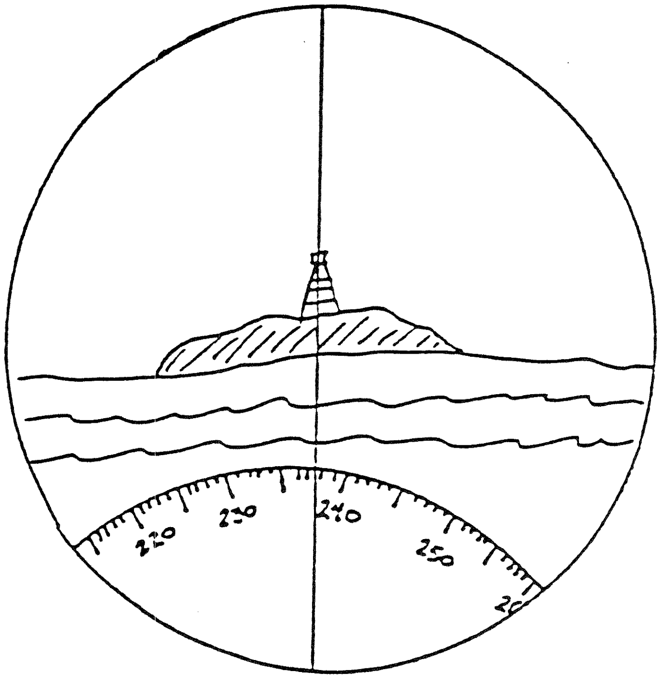

FIG. 8.2. View of a landmark with gyrocompass scale superimposed.

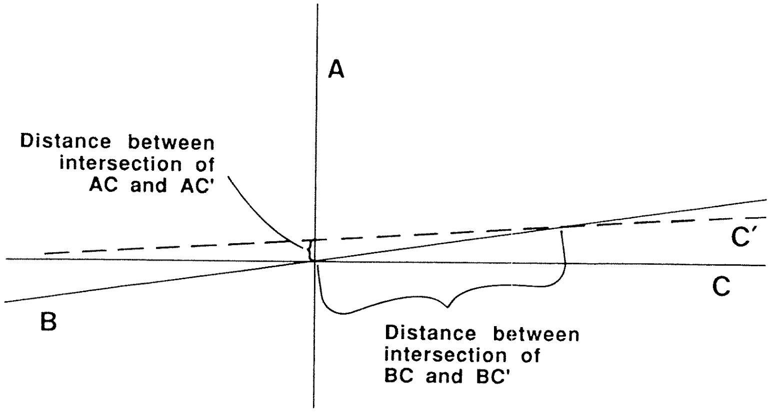

FIG. 8.3. Effects of errors of different sizes on the accuracy of the fix.

intersections of the lines of position are steep, then small errors in the observations themselves will have small effects on the locations of the intersections (see Fig. 8.3).

The bearing timer-recorder uses a wristwatch to time the fix intervals, and about 10 seconds before the next fix time, calls out "Standby to mark." This alerts the bearing takers that they should find their landmarks and aim their telescopic sights at them.

In a room just aft of the pilothouse sits another member of the team called the fathometer operator (position 6 in Fig. 8.1). This individual uses an instrument called an echo sounding depth finder or fathometer and is connected on the phone circuit with the bearing takers and the bearing timer-recorder. When the timer-recorder gives the "Standby to mark" signal, the fathometer operator reads the depth of water under the ship from the fathometer display and reports it to the bearing timer-recorder. This is recorded in the bearing log, and is later compared with the depth of water indicated on the chart at the plotted position. This comparison provides an additional check on the quality of the position fix.

At the time chosen for the fix observations, the bearing timer-recorder says "Mark," and the bearing takers observe and report the bearings of the landmarks they have been assigned. Thus the cycle begins again.

TECHNOLOGIES OF REPRESENTATION

How shall we think about the roles of technology and the people in this task performance? We are accustomed to thinking about technological systems as amplifiers of information-processing abilities or as intelligent intermediaries or agents who are also involved in the task performance. These are common metaphors for information-processing system design and they carry with them ways of thinking about the relations of people to technology that may prevent us from seeing the full range of possibilities (Hill, 1988; Hutchins, 1988). Looking in detail at this task has led me to think about the role of technology in cooperative work in a different way.

Problem Solving as Re-representation

In his seminal book, *The Sciences of the Artificial,* Herbert Simon (1981) said "Solving a problem simply means representing it so as to make the solution transparent" (p. 153). Of course, the meaning of "transparent" depends on the properties of the processor that must interpret the representation. Still, it is a powerful point, and although Simon had theorem proving in mind when he made it, it is also true of many other kinds of problem solving. The practice of moving a problem from one domain to another and solving it there has been an important technique in mathematics ever since Descartes' introduction of coordinates, a domain into which many problems have been moved and in which many solutions become apparent.[6] Gauss proved that a 17-sided regular polygon can be constructed with a straightedge and compass (the last previous regular polygon construction proof having come 2,000 years earlier), by using the knowledge that quadratic equations can be solved with straightedge and compass and showing that the construction of the 17-sided polygon was tantamount to solving a series of quadratic equations (Stewart, 1977). As useful as Cartesian coordinates are, many problems in relative motion that seem difficult in Cartesian coordinates become trivial when expressed in polar coordinates. Many instances of "Aha!" insight occur when a problem expressed in one way is re-represented in another in which the answer is literally staring the problem solver in the face (Gardner, 1978).

Four Ways To Do Distance/Rate/Time Problems

To get a better feel for how the way a problem is represented can change what is required of the problem solver, let us consider just one computation in the fix cycle. Suppose the plotter has just plotted a fix and needs to compute the ship's speed based on the distance the ship has moved in the interval of time that elapsed between the current fix and the previous one.

In particular, suppose the two fix positions are 1,500 yards apart and three minutes have elapsed between the fix observations.

---

[6] Modern navigation is still exploiting Descartes' insight. The chart itself is a carefully constructed coordinate space in which the solutions to many navigation problems become apparent.

Now, we are going to consider the cognition required of a task performer doing this task under four different conditions. Each of these is a real possibility.

> Condition 1: The task performer has the following resources: paper and pencil, knowledge of algebra, knowledge of arithmetic, knowledge that there are 2,000 yards in a nautical mile and 60 minutes in an hour, and knowledge of the equation D = RT.
>
> Condition 2: The task performer has the same resources as in condition 1, except that instead of a paper and pencil, the task performer has a four-function pocket calculator.
>
> Condition 3: The task performer has either a three-scale nomogram of the sort shown in Fig. 8.4, or a nautical slide-rule of the sort shown in Fig. 8.5 and the knowledge required to operate whichever tool is present.
>
> Condition 4: The task performer has no material implements at all, but knows how to use what navigators call the "three-minute" rule.

It is impossible to specify in advance just how any particular person will actually do this task in any of these conditions, but if the person actually uses the resources in the ways they are intended to be used, it is not difficult to determine what is likely to be involved.

In Condition 1, the task performer will first have to use the knowledge of algebra to manipulate the formula $D = RT$ to the form $R = D/T$ so that rate can be solved for directly from the given values of D and T. Then, the distance in yards will have to be converted to the equivalent number of miles using the knowledge of the number of yards in a mile and the knowledge of arithmetic. The time in minutes will have to be converted to the equivalent number of hours using the knowledge of the number of minutes in an hour and, again, arithmetic. The distance measure must be divided by the time measure, arithmetic again, to get the rate. Of course, these things can be done in a different order; for example, the division could come before either of the unit conversions, or between them, but in any case all these things must be done in order to solve the problem.

The reader may want to try it as an exercise just to get a feel for the sort of work involved. I estimate that this problem would tax the abilities of many navy navigation practitioners. Not because the arithmetic is difficult, but because it is necessary to figure out what to do and how the various things that are done fit together to produce the desired solution. One may be perfectly capable of doing every one of the component subtasks in this problem, but fail completely for lack of ability to organize and coordinate the various parts of the solution with each other.

In the calculator version, the procedures for doing the arithmetic opera-

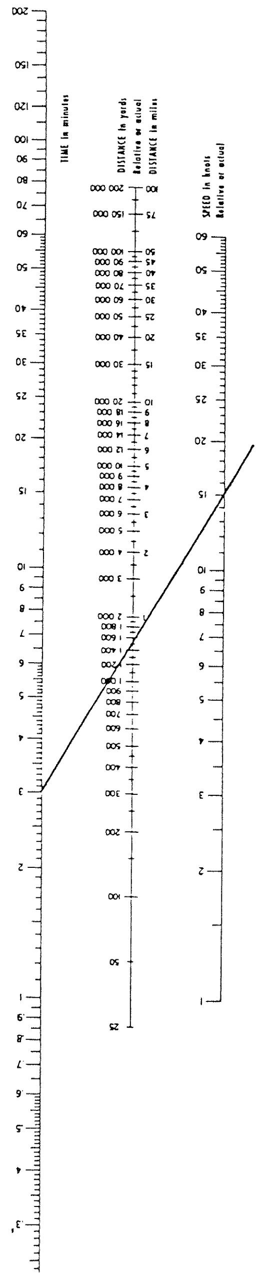

FIG. 8.4. Three-scale nomogram.

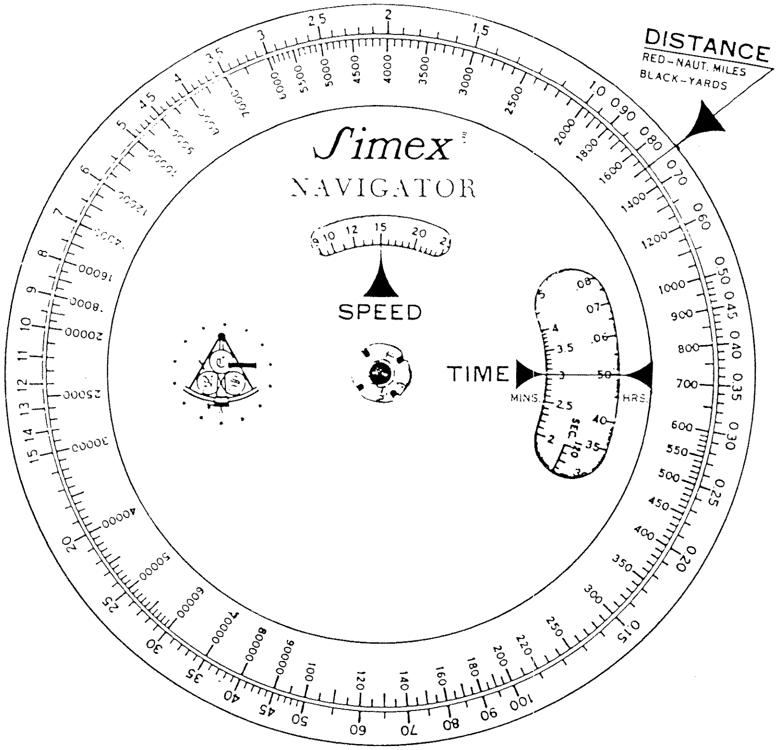

FIG. 8.5. Nautical slide rule.

tions of division and multiplication are restructured so that instead of constructing a pattern of symbols on a piece of paper and decomposing the problem to a set of operations on single-digit arithmetic arguments, values are keyed into the calculator and operator buttons are pushed. Also, depending on the order in which the steps are taken, it may be necessary to remember a previous result and enter it into a later operation after other operations have intervened. I think that this version of the task would also tax the abilities of many navigation practitioners, because the hard part is not doing the arithmetic but in deciding how to coordinate the arithmetic operations with each other. These tools give no support for that part of the task.

The paper and pencil condition and the calculator condition are alike in that they utilize completely general computational engines. The knowledge of the equation for distance, rate, and time and the knowledge of the constants required for the unit conversions are specific to the task, but they provide little help in structuring the actions of the task performer. As a

result, the procedures for doing the computation are complex. When we write them out at even the shallow level of detail previously given, we find that they contain many steps. If we actually got down to counting each symbol written on the paper or each key press on the calculator as a step (not an unusually detailed level for a cognitive analysis), we would find that they each run to many tens of steps.

Now consider the cognition required of the task performer in Condition 3. To use the nomogram, one finds the value of the time on the time scale and makes a mark there. One finds the value of the distance on the distance scale and makes a mark there. Then one draws a line through those two marks with a straightedge and reads the value of speed, in the desired units, where the line drawn intersects the speed scale. Now you may complain that the fact that these scales are already constructed in terms of the units set by the problem gives this condition an unfair advantage over the first two conditions. But that is part of the point. This is a very frequently occurring problem, and the nomogram is a tool designed specifically to make its solution easy. The use of the nautical slide-rule is very similar. It, like the nomogram, is a medium in which multiplication and division are represented as alignments of logarithmic scales. One aligns the distance index with the desired distance on the distance scale (could be yards or miles, both are represented side by side), aligns the elapsed time index with the desired time on the time scale (either minutes or hours, both are present side by side), and having done that the speed index will point to the speed in knots on the speed scale.

Having the scales in the units as set by the problem is helpful because it eliminates the need to convert one kind of unit into another, but it is more important to note that the knowledge of algebra is not required for this condition of the task. The nomogram and slide-rule transform the task from one of computation planning—figuring out what to divide by what—to one of simple manipulation of external devices. In the first two conditions, all that stands between the task performer and the nonsensical expressions $R=DT$ and $R=T/D$ is a knowledge of the syntax of algebraic transformations. When using the nomogram or the slide-rule, the structure of the artifacts themselves obviate or lock out such relations among the terms. The relations $D=RT$, $R=D/T$, and $T=D/R$ are built into the structure of the nomogram and slide-rule.[7]

The task performer still needs to know something, but the knowledge

---

[7] Looking at the nomogram, we see that the time and speed scales flank the distance scale on either side of it. A line drawn between any point on the speed scale and any point on the time scale intersects the distance scale at a point that is the averaged sum of the logarithm of the time and the logarithm of the speed. Sums of logarithms are products, so the very construction of the nomogram constrains the relationships among the terms to be of the correct type. Similarly, the slide-rule is constructed such that the distance reading is the angular sum of the logarithm of the speed and the logarithm of the time.

that is invoked to solve the problem with these tools is less complicated, and less general, than the knowledge required with the paper and pencil or calculator versions. A good deal of what needs to be done can be inferred from the structure of the artifacts themselves. They constrain the organization of action of the task performer by completely eliminating the possibility of certain syntactically incorrect relationships among the terms of the computation. Compared with Conditions 1 and 2, one may be reluctant to say that the answer was actually computed by the task performer in Condition 3. It seems that much of the computation was done by the tool used by the performer, or by the designer of that tool. The person somehow could get by doing less because the tool did more. But before we go that far, let us consider the task in Condition 4.

Where Condition 3 utilized specialized external artifacts, Condition 4 utilizes a specialized internal artifact. Since three minutes is one-twentieth of an hour and 100 yards is one-twentieth of a mile, the number of hundreds of yards (twentieths of a mile) a ship travels in three minutes (a twentieth of an hour) is its speed in nautical miles per hour. Thus, a ship that travels 1,500 yards in three minutes has a speed of 15 nautical miles per hour. In order to "see" the answer to the problem posed, the navigator need only imagine the number that represents the distance travelled in yards, 1500, with the last two digits removed: 15. The representation in which the answer is obvious is simply one in which the images of the characters that make up the numbers are manipulated.

Now the reader may really cry "foul play." "It's a special example with ginned-up numbers that permits this trick to be used," you say. In fact, this is not an unusual problem for a navigator. The most common interfix interval is three minutes precisely because this rule is so simple and so easy to use. The navigation team is capable of performing the fix cycle on two-minute, or even one-minute, intervals, but three minutes is more common—not because it meets the needs of the ship better than the other intervals, but because it meets them well enough, and it makes this computation so convenient.[8]

What are these tools contributing to the computations? It has now become commonplace to speak of technology, especially information-processing technology, as an amplifier of cognitive abilities. Cole and Griffin (1980) showed, however, that the appearance of amplification is an artifact of a commonly assumed perspective. When we concentrate on the product of the cognitive work, cultural technologies, from writing and mathematics to the kinds of tools we have considered here, appear to amplify the cognitive powers of their users. Using these tools, people can certainly do things they could not do without them. When we shift our focus to the process by which cognitive work is accomplished, however, we see something else. Every complex cognitive performance requires the application of some number of component cognitive abilities. Computing speed from distance and time with the calculator involves many component subtasks; remembering a symbolic expression, transforming the expression, determining which quantities correspond to which terms of the expression, mapping the expression to operations on the calculator, finding particular calculator keys, pressing the keys, and so forth. The application of these abilities must be organized in the sense that the work done by each component ability must be coordinated with that done by others. If we now consider doing the same task with the nomogram or with the three minute rule, we see that a different set of abilities is enlisted in the task. None of the component cognitive abilities has been amplified by the use of any of the tools. Rather, each tool presents the task to the user as a different sort of cognitive problem requiring a different set of cognitive abilities or a different organization of the same set of abilities.

There are two important things to notice about the computational technology of the piloting task. First, the existence of these specialized tools and techniques is evidence of a lot of cultural elaboration directed toward avoiding the use of algebraic reasoning and arithmetic. In fact, there is more than I have presented here. The problem could also have been solved by looking up the speed in a table of distances, rates, and times. We can only surmise that these are things that people are not good at.

Moreover, these tools and techniques permit the task performer to avoid doing algebraic reasoning and arithmetic by replacing those activities with aligning indices with numbers on scales, or imagining numerical representations and making simple transformations. Rather than amplifying the cognitive abilities of the task performers, or acting as intelligent agents in interaction with them, these tools transform the task the person has to do by mapping it into a domain where the answer or the path to the solution is apparent.

Perhaps this should also give us a new sense of expert systems. Clearly, a good deal of the expertise in the system is in the artifacts, both the external implements and the internal strategies; not in the sense that the artifacts are themselves intelligent or expert agents, but because the act of getting into

---

[8] The nautical slide-rule and nomogram are normally only used when the ship is away from land and the fix intervals are much longer than three minutes. When the cycle is performed on the shorter intervals of one or two minutes, speed is normally computed by conversion to the three minute standard. For example, if the ship travels 800 yards in two minutes, it would travel 1200 yards in three minutes, so its speed is 12 knots. And finally, regardless of the speed of the ship, as long as both the speed and the fix interval are constant, there is no need to recompute ship's speed to project the ship's position for the next fix. The distance traveled during the next interval will be the same as that covered in the last interval, so it can simply be spanned with dividers and laid on the projected track line without even determining its actual distance.

coordination with the artifact constitutes an expert performance by the person. These tools permit the people using them to do the tasks that need to be done while doing the kinds of things the people are good at: recognizing patterns, modeling simple dynamics of the world, and manipulating objects in the environment (Rumelhart, Smolensky, McClelland, & Hinton, 1986). At this end of the technological spectrum at least, the computational power of the system composed of person and technology is not determined primarily by the information-processing capacity internal to the technological device, but by way the technology exploits the cognitive resources of the task performer.

Propagation of Representational State Across External Media

Let us return now to the execution of the whole fix cycle by the navigation team. In light of the preceding discussion, we can now see the activity of the team differently. Each technological system involved provides a representation of the information about the relation of the ship to the world. When bearing takers align the hairline in the alidade sight with a landmark on the shore, they have imposed a representational state on the alidade in which the hairline also crosses the gyrocompass scale at a particular point. When they read that bearing and speak it into the phone, they have moved the representation of the bearing from the domain of the gyroscope scale to the domain of spoken words. When bearing timer-recorders write the bearing in the book, they impose a new representational state on the bearing book. The information is moved to the domain of written numbers. When plotters read the bearing and align the index of the plotting tool with its scale, they move the information to still another domain. And, when they plot the bearing, they move the information to a domain in which it can be integrated with other information of the same type to form a position fix. That is, drawing the line of position imposes a representational state on the chart. It is still a representation of the bearing of the landmark, but now it is in a medium in which its relationship to other lines of position establishes a position of the ship.

The chart is the domain in navigation where the representation of information about the relation of the ship to its surroundings becomes interpretable. It is the domain where it is easy to "see" the answer to the question, "Where is the ship?" The written entries in the bearing log provide a complete specification of the position of the ship with respect to the landmarks on the shore, but it is not easy to see the answer in those entries, nor is it easy to perform the subsequent computations required by the task in the representation as it appears in the bearing log. The task of the navigation team, therefore, is to propagate information about the directional relationships between the ship and known landmarks across a set of technological systems until it is represented on the chart. Between the situation of the ship in the world and the plotted position on the chart lies a bridge of technological devices. Each device (alidade, phone circuit, bearing log, etc.) supports a representational state, and each state is a transformation of the previous one. Each transformation is a trivial task for the person who performs it, but, placed in the proper order, these trivial transformations constitute the computation of ship's position.

THE COORDINATION OF HUMAN ACTIVITY

In the previous section, we saw the work of the members of the navigation team as the propagation of representational state across a number of media until it arrived on the chart. To accomplish this, the various activities that go into the navigation task must be coordinated with each other. At the beginning of the chapter, I noted that the fix cycle is executed by an individual working alone when the ship is not in restricted waters. This fact provides us with an opportunity to contrast the means of solving the coordination problem in the two work configurations.

The job of the solo watch stander is more difficult than the task facing any member of the navigation team because the lone task performer must not only operate all of the devices, but must also coordinate the use of each device with the uses of the others. The problems of coordination in the solo performance concern the control of the sequence of actions required to do the job. Recall that controlling the sequence of steps was just the problem faced by the navigator trying to compute the ship's speed from distance and time using paper and pencil or a calculator. In that case, the problem of sequence control was eliminated by the design of specialized tools on which the task is performed in a single step. For the lone watch stander doing the entire piloting task, the sequential control problem cannot be eliminated, but it is solved, as in many other military tasks, by a different sort of artifact, a procedure. A procedure is a plan for sequential action, and the task performer is expected to learn the procedure and use it as a guide in organizing his actions.

Coordination of Action in the Team Performance

When the navigation task is performed by the team, the coordination among the actions of the members of the team is not achieved by following a master procedure. Instead it emerges from the interactions among the members of the team.

Consider the bearing taker. This person coordinates his or her activity with the (timing) behavior of the timer-recorder. Upon hearing "Standby to mark" the bearing taker finds the landmark, and then waits for the "mark" signal. Upon hearing that, he or she reads and reports the bearing of the designated landmark. The bearing taker has thus delegated the control of some aspect of his or her own behavior to the timer-recorder. This person has delegated the control over some other aspects of his or her behavior to the device with which he or she interacts. The bearing taker reads the value that appears where the hairline crosses the gyrocompass scale. His or her behavior is nicely constrained by the two coordination activities. The bearing timer-recorder can only attend to one bearing report at a time, thus the two bearing takers must coordinate their reports. If one hears that another is already in the process of reporting a bearing, the first person must wait until the phone circuit is clear to make a report.

When a bearing taker reports a bearing, the timer-recorder coordinates his or her (recording) behavior with the bearing taker. Other activities are on hold while the timer-recorder attends to and records (perhaps doing both simultaneously) the bearing. The plotter waits for the first bearing to appear in the bearing log book and plots it while the other are being reported. By the time the plotter is ready for the second bearing, it is usually already recorded in the log, so it can be read and plotted.

In the team performance configuration, in the place of an executive we find an interrelated set of functional units. Each team member does a part of the job only when certain conditions appear in the task environment. Coordination among the activities of the team members arises because some of the conditions for each team member's actions are produced by the activities of the other members of the team.[9]

The Metronome of Execution

As we saw earlier, taking fixes at regular intervals greatly simplifies the computations involved in projecting the position of the ship at the next fix time. For this reason and others, the initiation of the fix cycle is carefully timed. This results in a remarkable periodicity in the performance of this task.

The coordination of the whole system with the meter of time is accomplished when the timer-recorder coordinates with a wristwatch and the others coordinate with the timer-recorder. The timer-recorder's coordination with the watch requires a maintained (a) vigilance to the watch, and (b) a test of when it is time to take another round. For the latter, the time-recorder must have a procedure for determining when the next round should fall and a way of determining when that time has been reached. Lack of vigilance, owing to the appropriation of attention by other tasks, may cause the timer to miss a mark. The plotter, who shares the physical environment of the timer apparently sometimes participates in that task redundantly and has been observed to comment, "Isn't it about time for a round?"

Coordination by Mutual Constraint

Notice that in the group performance mode, the sequence of actions to be taken need not be explicitly represented anywhere in the system. If participants know how to coordinate their activities with the technologies and people with which they interact, the global structure of the task performance will emerge from the local interactions of the members. The structure of the activities of the group is determined by a set of local computations rather than by the implementation of the sort of global plan that appears in the solo performer's procedure. In the team situation, a set of behavioral dependencies are set up. These dependencies shape the behavior pattern of the group. Similar effects have been observed by Kraut, Galegher, and Egido (1987-88) in their study of the distribution of labor among scientific collaborators.

Nominal and Real Divisions of Labor

In ideal conditions, the nominal division of labor depicted in the preceding sections is a reasonable description of what people do. But such conditions rarely prevail. Most of the time, there are small problems being encountered and solved, small errors being committed and corrected, and little bits of interaction structure being broken and repaired. In these more usual conditions, the nominal division of labor is routinely violated. Let me illustrate this with three brief examples.

Sometimes, bearing takers are not able to find the landmarks they have been assigned to observe. In the following exchange, the starboard bearing taker needs additional information to resolve an ambiguity. Here, SW is the starboard wing bearing taker and B is a qualified watch stander working as bearing timer-recorder.

SW: (Is it)The one on the left or the one on the right?
B: The one on the left, O.K.?
SW: Yah, I got it.

---

[9] This aspect of the coordination of the activities of the team could be well modeled by a production system in which each team member would be modeled as a set of productions and the working memory of the system would model the environment to which they respond and in which they produce states for each other.

In this case, the bearing timer-recorder is called on to do part of the bearing taker's job of identifying the landmark. When the confusion or lack of knowledge is more profound, it is simply impossible to communicate enough information over the phone circuit, and someone has to go in person to the wing and do more of the bearing taker's job. A little later in the same exit from port, the starboard bearing taker was unable to find the north end of the 10th Avenue terminal. The plotter P, who is also the most qualified and highest ranking member of the team, went onto the wing to point it out. On the wing, P put his arm over SW's shoulders and aimed his body in the right direction.

> P: The north one, all the way up.
>
> SW: O.K.
>
> P: If you can't see the light, just shoot the tangent right on the tit of the, the last end of the pier there.
>
> SW: O.K., that pier, where those two . . .
>
> P: Yah, all the way at the end.
>
> SW: Alright.
>
> P: There should be a light out there but if you can't see the light out there at the end of the pier (when we get in position), just shoot the end of the pier.

In the previous example, it was the bearing timer-recorder who temporarily left his duties to take on those of the bearing taker. This time it is the plotter.

At one point during an entry to port, the plotter was called away from the chart table for a consultation with the ship's captain just at the beginning of a fix cycle. When the bearings had been reported, the bearing timer-recorder reached over and set the index of the plotting tool to the first reported bearing. Upon returning a few seconds later, the plotter found that someone had done the first step of the plotting job.

This type of helping action occurs frequently. If they did not, if each individual was required to produce flawless performance entirely alone, the system would grind to a halt everytime any member of the team was unable to do his or her part of the job. Not only are members of the team responsible for their own jobs, they seem also to take responsibility for all parts of the process to which they can contribute.

The discussion of the coordination of action showed the people to be a sort of connecting tissue that holds the hardware of the technological systems together. The devices do not communicate with each other, it is the people who move the information from one device to another in performing the task of navigation.[10] Here we have seen that the flexibility and robustness of the system stems from the ability of this connecting tissue to adapt to changing circumstances. Sometimes the changing circumstances are consequences of problems in the connecting tissue itself (e.g., when a team member is unable to do his job), and sometimes it is a consequence of problems with the hard technology. In all cases, the management of the deployment of human resources, the on-line negotiation of the distribution of labor, is essential to the operation of the system. If the distribution of labor was fixed, the system would surely fail whenever one of its components failed.

The distribution of labor can only be negotiated if the distribution of knowledge and ability is at least partially redundant. In the organizational literature, this has been called "redundancy of function." Emery and Trist (1973) argued that such redundancy is necessary for the operation of adaptive, self-regulating systems in variable environments. Morgan (1986) called such systems "holographic" because "capacities relevant for the functioning of the whole are built into the parts" (p. 99). These notions raise the issue of the need to negotiate the division of labor, but do not address the issues of the way that negotiation proceeds or the conditions under which it can work. Emery and Trist (1973) and, later, Stratton and Flynn (1980) saw values as the coordinating mechanism among functionally redundant units. As noted earlier, members of the team seem to take responsibility for doing whatever they know can be done to ensure the timely propagation of information through the system. But values can only provide the disposition to act. To actually help, one must be aware of the need to help.

PRODUCTION AND REPRODUCTION IN COOPERATIVE WORK

The movement of people through careers or the fact of mortality more generally forces human systems to solve two problems at once. They must both produce whatever they produce, and, if they are to endure, they must reproduce themselves at the same time. The need to reproduce the system itself involves parts of what McGrath (chap. 2 in this volume) called the "support function," the system's contribution to its component parts, and the "well being function," the system's contribution to its own viability. As

---

[10] As an antidote to rampant anthropomorphisms in artificial intelligence, Michael Cole often stressed the idea that people do not communicate with machines, rather their communication with other people can be mediated by machines. In this setting perhaps we could say that the devices do not communicate with each other, except when their communication with each other is mediated by people.

people leave the system, replacement expertise must somehow be created. Sometimes this reproductive function is delegated to special institutions like schools, but much of what needs to be known is, and perhaps can only be, acquired in the work place.[11]

An Overlapping Distribution of Knowledge

The flexibility and robustness of the system described in the previous section is made possible by the overlapping distribution of knowledge and ability across the members of the team. This distribution of knowledge that is so important to the synchronic operation of the team is a consequence of the way the system replicates itself.

Knowledge in cooperative tasks is frequently assumed by analysts to be partitioned among individuals in an exhaustive and mutually exclusive manner such that the sum of the individuals' knowledge is equal to the total required, and there is little or no overlap. Consider the knowledge required to perform just the input portion of the basic fix cycle. This requires the knowledge of the bearing takers, the bearing timer-recorder, and the plotter. We could imagine designing an experiment along these lines by training up a person to perform each of these roles and then putting them in interaction with each other. This assumes no history for the participants except that they are each trained to do a particular job. This would result in a distribution of knowledge as shown in Fig. 8.6. Here the knowledge required to do each job is represented by a region in the pie.

It is certainly possible to organize a functional system along these lines, but in fact, outside of experimental settings, this is a very rare knowledge distribution pattern because it is very vulnerable to breakdown. If, for any reason, one of the members of the team is unable to perform, the whole system will fail. At the other end of the knowledge distribution spectrum, one can imagine a system in which everyone knows everything about the task. This is the pattern called for by Emery and Trist (1973) and by Morgan's (1986) "holographic" conception, but it too is a rare pattern, because if acquiring the knowledge costs anything at all, this pattern will be very expensive to produce. More commonly, there is substantial sharing of knowledge between individuals with the task knowledge of more expert performers completely subsuming the knowledge of those who are less experienced. Splitting the task into coordinated fragments permits relatively less skilled people to contribute to task performance.

In many human systems, as people become more skilled they move on to other roles in the task performance group, making way for less skilled people behind them and replacing the more expert people before them who advance or leave the system. This is what we observe in the case of the development of navigation skills aboard ship. The bearing taker knows how to do that job, but because of the interaction with the bearing timer-recorder, the person also knows something about the timer-recorder's job (see Fig. 8.7a). The bearing timer-recorder knows how to do that job, but also knows all about being a bearing taker, having once held that job. Furthermore, the timer-recorder knows a good deal about the activities of the plotter because they share the chart table. What the bearing timer-recorder knows is shown in Fig. 8.7b. Finally, a competent plotter knows how to plot, but also knows everything the bearing timer-recorder and bearing takers know having worked both those jobs before advancing to plotting. What the plotter knows is shown in Fig. 8.7c. The distribution of knowledge that is the sum of these individual expertises is shown in Fig. 8.7d. Thus, this movement through the system with increasing expertise results in a pattern of overlapping expertise, with knowledge of the entry-level tasks most redundantly represented and knowledge of expert-level tasks least redundantly represented.

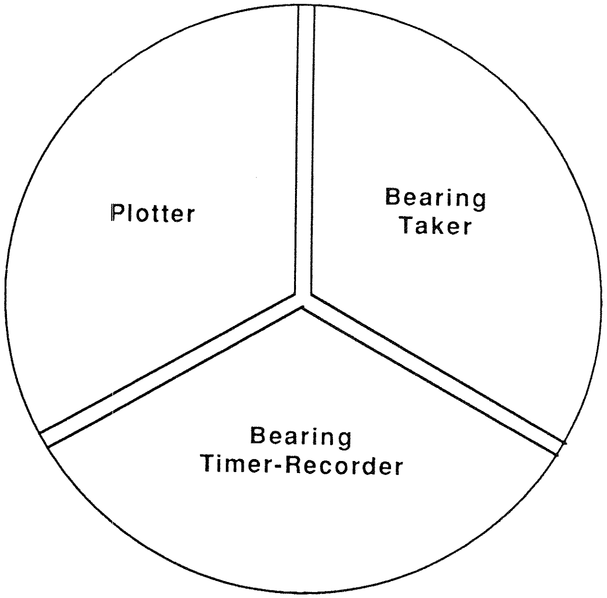

FIG. 8.6. Distribution of knowledge among individuals assuming no interaction and no history of experience in each other's jobs.

---

[11] Cicourel, (chap. 9 in this volume) discusses aspects of medical knowledge that can only be acquired in the practice of medicine in a social setting characterized by a distribution of expertise similar to the one described here.

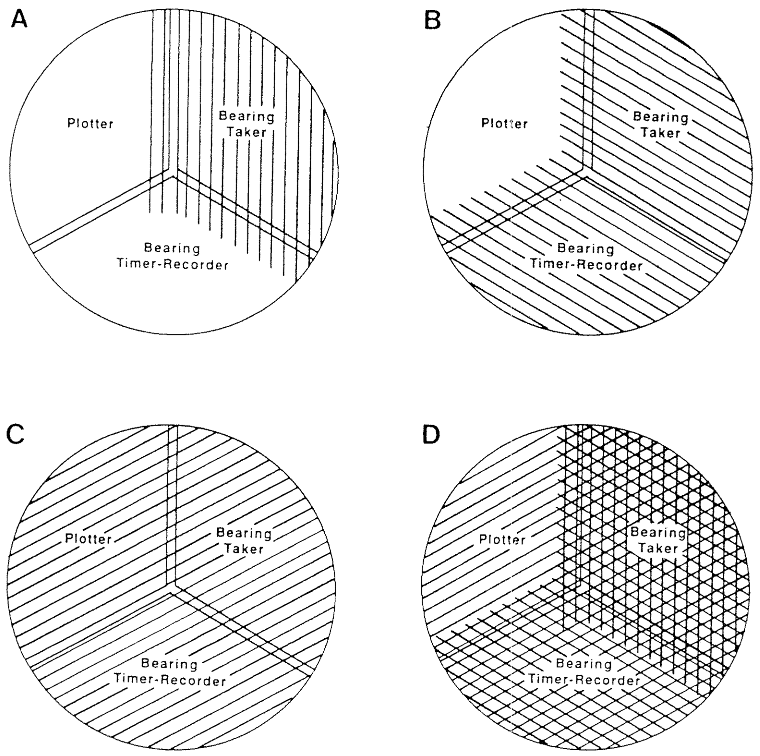

FIG. 8.7. Overlap in the distribution of knowledge as a result of interaction between individuals and experience in each other's jobs.

The Horizon of Observation

Let us refer to the outer boundary of the portion of the task that can be seen or heard by each team member as that person's horizon of observation. The scope of the horizons of observation of the members of a work group is important for two reasons. First, lines of communication and limits on observation of the activities of others have consequences for the knowledge acquisition process. This is so because they define the portion of the task environment available as a learning context to each task performer. As Hackman (1987) argued, making some parts of a joint task more visible to other participants permits group members to educate each other. Second, as already noted, the awareness of the need for help is often a prerequisite for help. The extent to which a system can benefit from the functional redundancy of its parts will depend on the extent to which the parts are aware of needs for redundant functioning. Technology has a key role here because the horizons of observation of the members of the work group are often defined by technology.

Open Interactions

The physical arrangement of tools and work stations together with the local ethos of work can affect the horizon of observation, and thus the opportunities to learn, of members of the team. Here is an example of a learning event that was enabled by the fact that the plotter, the bearing timer-recorder, and the keeper of the deck log all share the same work space.

On a previous at sea period, D, the deck log keeper had served as bearing timer-recorder, but D's performance there was less than satisfactory. That is the job that was next in line for him, however, and D was anxious to acquire the skills required to perform the job. One of the most important aspects of the bearing timer-recorder's job is knowing when particular landmarks will be visible to the bearing takers on the wings. One complication of this judgment is the fact that a large convex mirror is mounted outside the pilot house windows just in front of the port wing bearing taker's station. The mirror is there so that the commanding officer, who sits inside the pilot house, can see the part of the flight deck that lies aft of the pilot house. Unfortunately, the mirror obstructs the port bearing taker's view forward and the bearing timer-recorder must be able to judge from his position at the chart table whether or not the port wing bearing taker's view of a chosen landmark will be blocked by the mirror.

The plotter, P, the bearing timer-recorder, B, and D, were all standing at the chart table. The ship had just entered the mouth of the harbor and the team was running the fix cycle on two-minute intervals. The previous fix taken at 36 minutes after the hour, called "time 36," was complete, and P had just finished plotting the dead reckoned track out through times 38 and 40. B indirectly solicited P's assistance in deciding which landmarks should be shot for the next round of bearings. D stood by, watching what B and P were doing. All of the pointing they did in this interchange was to the chart itself.

1. B: Last set still good? O.K. Ballast Point, light Zulu.
2. P: Here's (time) 36 [pointing to the DR position on the chart]
3. B: So it would be that [pointing to light Zulu], that [pointing to Bravo Pier] . . .
4. P: One, two, three. Same three. Ballast Point, Bravo. And the next one . . .
5. B: (Time) 40 should be, Ballast Point . . .

6. P: Front Range, Bravo.
7. B: And Balla . . .
8. D: He may not be able to see Front Range.
9. B: Yah.
10. P: Yah, he can. Once we get up here [pointing to the ship's projected position for the next fix].
11. B: Yah. Up there O.K.
12. P: Down here [pointing to ships current position] he can't. It's back of the mirror, but as you come in it gets enough so that you can see it.

Because what B and P are doing is within D's horizon of observation, D has a chance to see how the landmarks are chosen. Furthermore, the fact that the decision about which landmarks to shoot is made in an interaction opens the process to D in a way that would not be the case if a single person were making the decision alone. In utterance 8, D raises the possibility that the port wing bearing taker may not be able to see the landmark. Three days earlier, on another sea and anchor detail, D had made the same suggestion about the mirror blocking the port wing bearing taker's view and P had agreed. In the present circumstances, however, D's caveat is inappropriate. B and P have already anticipated the problem raised by D, and they jointly counter D's objection, each building on what the other has said. Clearly, if D did not share the work space with B and P or if there was a strict division of labor such that people did not monitor and participate in the actions of their fellows, this opportunity for D to have even peripheral involvement in that task that will someday be his would be lost. Furthermore, D's horizon of observation is extended because the decision making about landmarks is conducted as an interaction between B and P.

Open Tools

However, being in the presence of others who are working is not always enough by itself. In the previous example, we saw that the fact that the work was done in an interaction between members opened it to other members of the team. In a similar way, the design of tools can affect their suitability for joint use or for demonstration and may thereby constrain possibilities for knowledge acquisition. The interaction of a task performer with a tool may or may not be open to others depending on the nature of the tool itself. The design of a tool may change the horizon of observation for those in the vicinity of the tool. The navigation chart is an explicit graphical depiction of position and motion so it is easy to see certain aspects of solutions. The chart representation presents the relevant information in a form such that much of the work can be done by perceptual inferences. The work a chart does is performed on its surface—all at the device interface, as it were—but watching someone work with a chart is much more revealing of what is done to perform the task than watching someone work with a calculator or a computer.

The openness of a tool can also affect its use as an instrument in instruction. When the bearing timer-recorder chooses a set of landmarks that result in lines of position with shallow intersections, it is easy to show, on the chart, the consequences of such actions and the nature of the remedy required. Figure 8.8 shows a fix that resulted from landmark assignments made by the bearing timer-recorder. Bearings off to the side of the ship rather than ahead or astern are called *beam* bearings. After plotting this fix and observing how it came out, the plotter scolded the bearing timer-recorder.

> C: What did you take a bunch of beam bearings for? Why ain't you shooting up there [points out the front window of the bridge] some

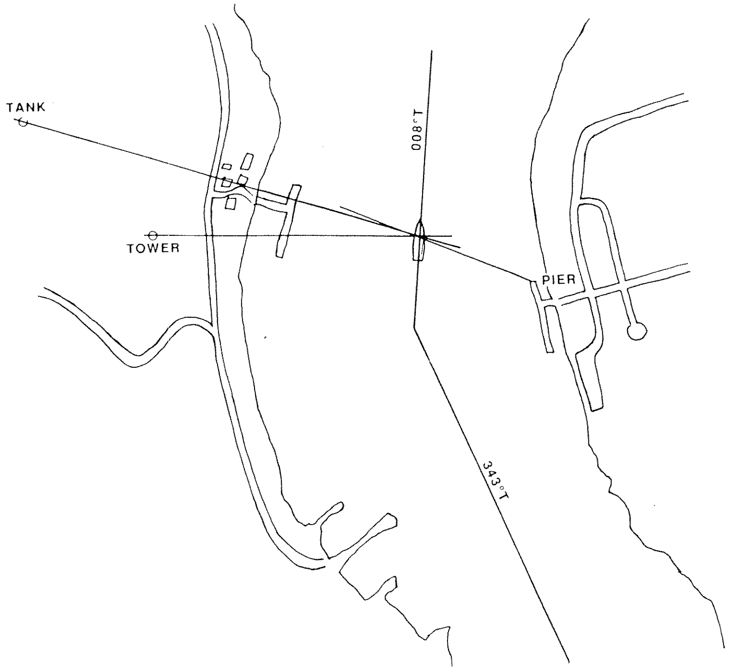

FIG. 8.8. Fix based on "beam" bearings resulting in shallow intersections of the lines of position.

> place? Look what you did! [points to the chart] You shot three beam bearings. You shot three beam bearings. You better tell 'em to shoot from up ahead some place.

Once the fix was plotted of course, it was easy for the bearing timer-recorder to see the nature of his error. Imagine how much more difficult it would be to explain the inadequacy of the landmark assignment if the lines of position were represented as equations to be punched into a calculator, rather than as lines drawn on the chart.

CONCLUSION

When we consider the activities of the navigation team in any particular navigation exercise, the costs associated with the added burdens of communication and coordination in the socially distributed form of this task could be seen, following Steiner (1972), as process losses for the system. Clearly, there is a good deal of inefficiency in the coordination of the activities of the various team members with each other. One can argue, even in the particular instance, that some of the losses are compensated for by the increased information-processing speed that comes with parallel task performance. After all, that is why the task is done by a team when the time pressures are great. Still, the costs are high. McGrath (chap. 2 in this volume) has argued, however, that we must take a broader perspective in which we consider not only the production function of group activity, but the support and well-being functions as well. Our concern should not be simply for the efficiency of any particular performance of the task, but for the quality of the performances of many tasks across a span of time on the scale of an individual team member's career. From this perspective, we see that some of what appear as losses in particular instances are in fact investments in future performances. Some of these costs are payment for the maintenance of a base of expertise that is periodically threatened by the departure of an experienced member of the team.

Let me now recapitulate the main points of the argument. I have tried to show that in the world of large ship piloting, the technological devices are better seen as media for representation than as amplifiers or surrogates for cognitive abilities. The central computations of the navigation tasks are accomplished by the propagation of information across representations and representational media. There is a nominal division of labor among the members of the team that is supposed to ensure the timely propagation of information, but the need to violate the mandated division of labor in order to avoid breakdowns arises frequently. The negotiation of the division of labor that permits the system to respond robustly to failures of individual team members requires an overlapping distribution of expertise; that is, some level of functional redundancy. In addition to satisfying its production function, the system must also reproduce itself, and its mode of reproduction, as reflected in the career trajectories of the individual group members, both replaces lost expertise and maintains the desired functional redundancy. The space through which the career trajectory passes, the jobs that the members do, and the opportunities to learn in those jobs are all defined in large part by the properties of the technology. Many tools in this setting are open tools in the sense that their use in public and observable in its details by other members of the team. These open tools provide many learning contexts.

The observations I have presented here should not be taken as either claims that the current ways of doing things are optimal, or as arguments against the automation of the tasks of ship navigation. The existing system of ship navigation is problematic in many respects and very probably could be improved by technological innovations. I have tried to show that the robustness of the system in the face of component failures is dependent on the openness or public nature of much of the work of the navigation task. For example, the detection of error requires access to errorful performance and the correction of error requires a functionally redundant distribution of knowledge. There is a danger because the public nature of work that creates these required conditions seems particularly vulnerable to the introduction of high-tech solutions to the narrowly conceived production problem. Although some promising work on computer support for collaborative work is now underway (see Lakin, chap. 17; and Landow, chap. 15 in this volume; and CSCW, 1986, 1988), the typical high-tech widget is still a single operator system where the interaction between user and device is a private world of activity. Rather than an argument against automation, take this to be a call for expanded design criteria.

The low technology system examined here incorporates ways for operators to learn their jobs while working, redundant means for the detection and the correction of error, and possibilities for adaptation and continued functioning in the face of failures of component parts. It would be unfortunate indeed if these clearly desirable features were overlooked and left behind in the pursuit of an efficient automated means of fostering productivity.

ACKNOWLEDGMENTS

The research on which this chapter is based was supported by the Independent Research program of the Navy Personnel Research and Development Center, San Diego, by the Personnel and Training Research programs, Psychological Sciences Division, Office of Naval Research, under Contract No. N00014-85-C-0133, and by a fellowship from the John D. and Catherine T. MacArthur Foundation. The views and conclusions presented in this chapter are those of the author and should not necessarily be interpreted as representing the official policies, either expressed or implied, of the sponsoring agencies.

REFERENCES

Cole, M., & Griffin, P. (1980). Cultural amplifiers reconsidered. In D. R. Olson (Ed.), *The social foundations of language and thought* (pp. 343–364). New York: Norton.

CSCW. (1986). *Proceedings of the conference on computer supported cooperative work.* Austin, TX: Microelectronics & Computer Technology Corporation.

Emery, F. E., & Trist, E. L. (1973). *Towards a social ecology.* New York: Plenum.

Gardner, M. (1978). *Aha! insight.* New York: W. H. Freeman.

Hill, W. C. (March 1988). *The mind at AI: Horseless carriage, mousetrap, card trick, clock* (Internal report). Austin, TX: Microelectronics & Computer Technology Corporation.

Hackman, J. R. (1987). The design of work teams. In J. W. Lorsch (Ed.), *Handbook of organizational behavior* (pp. 315–342). Englewood Cliffs, NJ: Prentice-Hall.

Hutchins, E. (1988). Metaphors for interface design. In M. M. Taylor, F. Neel, & D. G. Bouwhuis (Eds.), *The structure of multimodal dialogue* (pp. 26–44). Amsterdam: North Holland.

Kraut, R. E., Galegher, J., & Egido, C. (1987–88). Relationships and tasks in scientific collaboration. *Human-Computer Interaction, 3,* 31–58.

Morgan, G. (1986). *Images of organization.* Beverly Hills, CA: Sage Publications.

Rumelhart, D., Smolensky, P., McClelland, J., & Hinton, G. (1986). Schemata and sequential processes. In J. McClelland, & D. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 7–57). Cambridge, MA: MIT Press.

Simon, H. (1981). *The sciences of the artificial* (2nd ed.). Cambridge, MA: MIT Press.

Steiner, I. D. (1972). *Group process and productivity.* New York: Academic Press.

Stratton, W. E., & Flynn, W. R. (1980). Ecological theory and organization development. In T. G. Cummings (Ed.), *Systems theory for organization development* (pp. 307–322). New York: Wiley.

Stewart, I. (1987). Gauss. In *Scientific genius and creativity: Readings from Scientific American* (pp. 40–49). New York: W. H. Freeman.