

The augmented knowledge workshop

by DOUGLAS C. ENGELBART, RICHARD W. WATSON, and JAMES C. NORTON

Stanford Research Institute
Menlo Park, California

CONCEPT OF THE KNOWLEDGE WORKSHOP

This paper discusses the theme of augmenting a knowledge workshop. The first part of the paper describes the concept and framework of the knowledge workshop. The second part describes aspects of a prototype knowledge workshop being developed within this framework.

The importance and implications of the idea of knowledge work have been described by Drucker.^{3,4} Considering knowledge to be the systematic organization of information and concepts, he defines the knowledge worker as the person who creates and applies knowledge to productive ends, in contrast to an "intellectual" for whom information and concepts may only have importance because they interest him, or to the manual worker who applies manual skills or brawn. In those two books Drucker brings out many significant facts and considerations highly relevant to the theme here, one among them (paraphrased below) being the accelerating rate at which knowledge and knowledge work are coming to dominate the working activity of our society:

In 1900 the majority and largest single group of Americans obtained their livelihood from the farm. By 1940 the largest single group was industrial workers, especially semiskilled machine operators. By 1960, the largest single group was professional, managerial, and technical—that is, knowledge workers. By 1975-80 this group will embrace the majority of Americans. The productivity of knowledge has already become the key to national productivity, competitive strength, and economic achievement, according to Drucker. It is knowledge, not land, raw materials, or capital, that has become the central factor in production.

In his provocative discussions, Drucker makes extensive use of such terms as "knowledge organizations," "knowledge technologies," and "knowledge societies." It seemed a highly appropriate extension for us to coin "knowledge workshop" for re-naming the area of our special interest: the place in which knowledge workers do their work. Knowledge workshops have existed for centuries, but our special concern is their systematic improvement, toward increased effectiveness of this new breed of craftsmen.

Workshop improvement involves systematic change not only in the tools that help handle and transform the materials, but in the customs, conventions, skills, procedures, working methods, organizational roles, training, etc., by which the workers and their organizations harness their tools, their skills, and their knowledge.

Over the past ten years, the explicit focus in the Augmentation Research Center (ARC) has been upon the effects and possibilities of new knowledge workshop tools based on the technology of computer timesharing and modern communications.¹⁸⁻⁴¹ Since we consider automating many human operations, what we are after could perhaps be termed "workshop automation." But the very great importance of aspects other than the new tools (i.e., conventions, methods, roles) makes us prefer the "augmentation" term that hopefully can remain "whole-scope." We want to keep tools in proper perspective within the total system that augments native human capacities toward effective action.^{1-3, 10, 16, 18, 24}

Development of more effective knowledge workshop technology will require talents and experience from many backgrounds: computer hardware and software, psychology, management science, information science, and operations research, to name a few. These must come together within the framework of a new discipline, focused on the systematic study of knowledge work and its workshop environments.

TWO WAYS IN WHICH AUGMENTED KNOWLEDGE WORKSHOPS ARE EVOLVING

Introduction

First, one can see a definite evolution of new workshop architecture in the trends of computer application systems. An "augmented workshop domain" will probably emerge because many special-purpose application systems are evolving by adding useful features outside their immediate special application area. As a result, many will tend to overlap in their general knowledge work supporting features.

Second, research and development is being directed toward augmenting a "Core" Knowledge Workshop domain. This application system development is aimed expressly at supporting basic functions of knowledge

work. An important characteristic of such systems is to interface usefully with specialized systems. This paper is oriented toward this second approach.

NATURAL EVOLUTION BY SCATTERED NUCLEI EXPANDING TOWARD A COMMON "KNOWLEDGE WORKSHOP" DOMAIN

Anderson and Coover¹⁵ point out that a decade or more of application-system evolution is bringing about the beginning of relatively rational user-oriented languages for the control interfaces of advanced applications software systems. What is interesting to note is that the functions provided by the "interface control" for the more advanced systems are coming to include editors and generalized file-management facilities, to make easier the preparation, execution, and management of the special-purpose tools of such systems.

It seems probable that special application-oriented systems (languages) will evolve steadily toward helping the user with such associated work as formulating models, documenting them, specifying the different trial runs, keeping track of intermediate results, annotating them and linking them back to the users' model(s), etc. When the results are produced by what were initially the core application programs (e.g., the statistical programs), he will want ways to integrate them into his working notes, illustrating, labeling, captioning, explaining and interpreting them. Eventually these notes will be shaped into memoranda and formal publications, to undergo dialogue and detailed study with and by others.¹⁵

Once a significant user-oriented system becomes established, with a steady growth of user clientele, there will be natural forces steadily increasing the effectiveness of the system services and steadily decreasing the cost per unit of service. And it will also be natural that the functional domain of an application system will steadily grow outward: "as long as the information must be in computer form anyway for an adjacent, computerized process, let's consider applying computer aid to Activity X also."

Because the boundary of the Application System has grown out to be "next to" Activity X, it has become relatively easy to consider extending the computerized-information domain a bit so that a new application process can support Activity X. After all, the equipment is already there, the users who perform Activity X are already oriented to use integrated computer aid, and generally the computer facilitation of Activity X will prove to have a beneficial effect on the productivity of the rest of the applications system.

This domain-spreading characteristic is less dependent upon the substantive work area a particular application system supports than it is upon the health and vitality of its development and application (the authors of Reference 15 have important things to say on these issues); however, it appears that continuing growth is bound to occur in

many special application domains, inevitably bringing about overlap in common application "sub-domains" (as seen from the center of any of these nuclei). These special subdomains include formulating, studying, keeping track of ideas, carrying on dialogue, publishing, negotiating, planning, coordinating, learning, coaching, looking up in the yellow pages to find someone who can do a special service, etc.

CONSIDERING THE CORE KNOWLEDGE WORKSHOP AS A SYSTEM DOMAIN IN ITS OWN RIGHT

A second approach to the evolution of a knowledge workshop is to recognize from the beginning the amount and importance of human activity constantly involved in the "core" domain of knowledge work—activity within which more specialized functions are embedded.

If you asked a particular knowledge worker (e.g., scientist, engineer, manager, or marketing specialist) what were the foundations of his livelihood, he would probably point to particular skills such as those involved in designing an electric circuit, forecasting a market based on various data, or managing work flow in a project. If you asked him what tools he needed to improve his effectiveness he would point to requirements for aids in designing circuits, analyzing his data, or scheduling the flow of work.

But, a record of how this person used his time, even if his work was highly specialized, would show that specialized work such as mentioned above, while vital to his effectiveness, probably occupied a small fraction of his time and effort.

The bulk of his time, for example, would probably be occupied by more general knowledge work: writing and planning or design document; carrying on dialogue with others in writing, in person, or on the telephone; studying documents; filing ideas or other material; formulating problem-solving approaches; coordinating work with others; and reporting results.

There would seem to be a promise of considerable payoff in establishing a healthy, applications oriented systems development activity within this common, "core" domain, meeting the special-application systems "coming the other way" and providing them with well-designed services at a natural system-to-system interface.

It will be much more efficient to develop this domain explicitly, by people oriented toward it, and hopefully with resources shared in a coordinated fashion. The alternative of semi-random growth promises problems such as:

- (1) Repetitive solutions for the same functional problems, each within the skewed perspective of a particular special-applications area for which these problems are peripheral issues,
- (2) Incompatibility between different application software systems in terms of their inputs and outputs,

- (3) Languages and other control conventions inconsistent or based on different principles from one system to another, creating unnecessary learning barriers or other discouragements to cross usage.

In summary, the two trends in the evolution of knowledge workshops described above are each valuable and are complementary. Experience and specific tools and techniques can and will be transferred between them.

There is a very extensive range of "core" workshop functions, common to a wide variety of knowledge work, and they factor into many levels and dimensions. In the sections to follow, we describe our developments, activities, and commitments from the expectation that there soon will be increased activity in this core knowledge workshop domain, and that it will be evolving "outward" to meet the other application systems "heading inward."

BASIC ASSUMPTIONS ABOUT AUGMENTED KNOWLEDGE WORKSHOPS EMBEDDED IN A COMPUTER NETWORK

The computer-based "tools" of a knowledge workshop will be provided in the environment of a computer network such as the ARPANET.^{7,8,14} For instance, the core functions will consist of a network of cooperating processors performing special functions such as editing, publishing, communication of documents and messages, data management, and so forth. Less commonly used but important functions might exist on a single machine. The total computer assisted workshop will be based on many geographically separate systems.

Once there is a "digital-packet transportation system," it becomes possible for the individual user to reach out through his interfacing processor(s) to access other people and other services scattered throughout a "community," and the "labor marketplace" where he transacts his knowledge work literally will not have to be affected by geographical location.²⁷

Specialty application systems will exist in the way that specialty shops and services now do—and for the same reasons. When it is easy to transport the material and negotiate the service transactions, one group of people will find that specialization can improve their cost/effectiveness, and that there is a large enough market within reach to support them. And in the network-coupled computer-resource marketplace, the specialty shops will grow—e.g., application systems specially tailored for particular types of analyses, or for checking through text for spelling errors, or for doing the text-graphic document typography in a special area of technical portrayal, and so on. There will be brokers, wholesalers, middle men, and retailers.

Coordinated set of user interface principles

There will be a common set of principles, over the many application areas, shaping user interface features

such as the language, control conventions, and methods for obtaining help and computer-aided training.

This characteristic has two main implications. One, it means that while each domain within the core workshop area or within a specialized application system may have a vocabulary unique to its area, this vocabulary will be used within language and control structures common throughout the workshop system. A user will learn to use additional functions by increasing vocabulary, not by having to learn separate "foreign" languages. Two, when in trouble, he will invoke help or tutorial functions in a standard way.

Grades of user proficiency

Even a once-in-a-while user with a minimum of learning will want to be able to get at least a few straightforward things done. In fact, even an expert user in one domain will be a novice in others that he uses infrequently. Attention to novice-oriented features is required.

But users also want and deserve the reward of increased proficiency and capability from improvements in their skills and knowledge, and in their conceptual orientation to the problem domain and to their workshop's system of tools, methods, conventions, etc. "Advanced vocabularies" in every special domain will be important and unavoidable.

A corollary feature is that workers in the rapidly evolving augmented workshops should continuously be involved with testing and training in order that their skills and knowledge may harness available tools and methodology most effectively.

Ease of communication between, and addition of, workshop domains

One cannot predict ahead of time which domains or application systems within the workshop will want to communicate in various sequences with which others, or what operations will be needed in the future. Thus, results must be easily communicated from one set of operations to another, and it should be easy to add or interface new domains to the workshop.

User programming capability

There will never be enough professional programmers and system developers to develop or interface all the tools that users may need for their work. Therefore, it must be possible, with various levels of ease, for users to add or interface new tools, and extend the language to meet their needs. They should be able to do this in a variety of programming languages with which they may have training, or in the basic user-level language of the workshop itself.

Availability of people support services

An augmented workshop will have more support services available than those provided by computer tools.

There will be many people support services as well: besides clerical support, there will be extensive and highly specialized professional services, e.g., document design and typography, data base design and administration, training, cataloging, retrieval formulation, etc. In fact, the marketplace for human services will become much more diverse and active.²⁷

Cost decreasing, capabilities increasing

The power and range of available capabilities will increase and costs will decrease. Modular software designs, where only the software tools needed at any given moment are linked into a person's run-time computer space, will cut system overhead for parts of the system not in use. Modularity in hardware will provide local configurations of terminals and miniprocessors tailored for economically fitting needs. It is obvious that cost of raw hardware components is plummeting; and the assumed large market for knowledge workshop support systems implies further help in bringing prices down.

The argument given earlier for the steady expansion of vital application systems to other domains remains valid for explaining why the capabilities of the workshop will increase. Further, increasing experience with the workshop will lead to improvements, as will the general trend in technology evolution.

Range of workstations and symbol representations

The range of workstations available to the user will increase in scope and capability. These workstations will support text with large, open-ended character sets, pictures, voice, mathematical notation, tables, numbers and other forms of knowledge representation. Even small portable hand-held consoles will be available.¹³

Careful development of methodology

As much care and attention will be given to the development, analysis, and evaluation of procedures and methodology for use of computer and people support services as to the development of the technological support services.

Changed roles and organizational structure

The widespread availability of workshop services will create the need for new organizational structures and roles.

SELECTED DESCRIPTION OF AUGMENTED WORKSHOP CAPABILITIES

Introduction

Within the framework described above, ARC is developing a prototype workshop system. Our system does not

meet all the requirements outlined previously, but it does have a powerful set of core capabilities and experience that leads us to believe that such goals can be achieved.

Within ARC we do as much work as possible using the range of online capabilities offered. We serve not only as researchers, but also as the subjects for the analysis and evaluation of the augmentation system that we have been developing.

Consequently, an important aspect of the augmentation work done within ARC is that the techniques being explored are implemented, studied, and evaluated with the advantage of intensive everyday usage. We call this research and development strategy "bootstrapping."

In our experience, complex man-machine systems can evolve only in a pragmatic mode, within real-work environments where there is an appropriate commitment to conscious, controlled, exploratory evolution within the general framework outlined earlier. The plans and commitments described later are a consistent extension of this pragmatic bootstrapping strategy.

To give the reader more of a flavor of some of the many dimensions and levels of the ARC workshop, four example areas are discussed below in more detail, following a quick description of our physical environment.

The first area consists of mechanisms for studying and browsing through NLS files as an example of one functional dimension that has been explored in some depth.

The second area consists of mechanisms for collaboration support—a subsystem domain important to many application areas.

The third and fourth areas, support for software engineers and the ARPANET Network Information Center (NIC), show example application domains based on functions in our workshop.

General physical environment

Our computer-based tools run on a Digital Equipment Corporation PDP-10 computer, operating with the Bolt, Beranek, and Newman TENEX timesharing system.⁹ The computer is connected via an Interface Message Processor (IMP) to the ARPANET.^{7,8} There is a good deal of interaction with Network researchers, and with Network technology, since we operate the ARPA Network Information Center (see below).³⁹

There is a range of terminals: twelve old, but serviceable, display consoles of our own design,²⁶ an IMLAC display, a dozen or so 30 ch/sec portable upper/lower case typewriter terminals, five magnetic tape-cassette storage units that can be used either online or offline, and a 96-character line printer. There are 125 million characters of online disk storage.

The display consoles are equipped with a typewriter-like keyboard, a five-finger keyset for one-handed character input, and a "mouse"—a device for controlling the position of a cursor (or pointer) on the display screen and for input of certain control commands. Test results on the mouse as a screen-

selection device have been reported in Reference 25, and good photographs and descriptions of the physical systems have appeared in References 20 and 21.

The core workshop software system and language, called NLS, provides many basic tools, of which a number will be mentioned below. It is our "core-workshop application system."

During the initial years of workshop development, application and analysis, the basic knowledge-work functions have centered around the composition, modification, and study of structured textual material.²⁶ Some of the capabilities in this area are described in detail in Reference 26, and are graphically shown in a movie available on loan!⁴¹—

The structured-text manipulation has been developed extensively because of its high payoff in the area of applications-system development to which we have applied our augmented workshop. We have delayed addition of graphic-manipulation capabilities because there were important areas associated with the text domain needing exploration and because of limitations in the display system and hardcopy print-out.

To build the picture of what our Core Knowledge Workshop is like, we first give several in-depth examples, and then list in the section on workshop utility service some "workshop subsystems" that we consider to be of considerable importance to general knowledge work.

STUDYING ONLINE DOCUMENTS

Introduction

The functions to be described form a set of controls for easily moving one around in an information space and allowing one to adjust the scope, format, and content of the information seen.^{26,41}

Given the addition of graphical, numerical, and vocal information, which are planned for addition to the workshop, one can visualize many additions to the concepts below. Even for strictly textual material there are yet many useful ideas to be explored.

View specifications

One may want an overview of a document in a table-of-contents like form on the screen. To facilitate this and other needs, NLS text files are hierarchically structured in a tree form with subordinate material at lower levels in the hierarchy.²⁶

The basic conceptual unit in NLS, at each node of the hierarchical file, is called a "statement" and is usually a paragraph, sentence, equation, or other unit that one wants to manipulate as a whole.

A statement can contain many characters—presently, up to 2000. Therefore, a statement can contain many lines of text. Two of the "view-specification" parameters—depth in the hierarchy, and lines per statement—can be controlled during study of a document to give various overviews of it. View specifications are given with highly abbreviated control codes, because they are used very frequently and their quick specification and execution make a great deal of difference in the facility with which one studies the material and keeps track of where he is.

Examples of other view specifications are those that control spacing between statements, and indentation for levels in the hierarchy, and determine whether the identifications associated with statements are to be displayed, which branch(es) in the tree are to be displayed, whether special filters are to be invoked to show only statements meeting specified content requirements or whether statements are to be transformed according to special rules programmed by the user.

Moving in information space

A related viewing problem is designating the particular location (node in a file hierarchy) to be at the top of the screen. The computer then creates a display of the information from that point according to the view specifications currently in effect.

The system contains a variety of appropriate commands to do this; they are called jump commands because they have the effect of "jumping" or moving one from place to place in the network of files available as a user's information space.^{26,33-39}

One can point at a particular statement on the screen and command the system to move on to various positions relative to the selected one, such as up or down in the hierarchical structure, to the next or preceding statement at the same hierarchical level, to the first or last statement at a given level, etc.

One can tell the system to move to a specifically named point or go to the next occurrence of a statement with a specific content.

Each time a jump or move is made, the option is offered of including any of the abbreviated view specifications—a very general, single operation is "jump to that location and display with this view."

As one moves about in a file one may want to quickly and easily return to a previous view of the path as one traverses through the file and the specific view at each point, and then allowing return movement to the most recent points saved.

Another important feature in studying or browsing in a document is being able to quickly move to other documents cited.

There is a convention (called a "link") for citing documents that allows the user to specify a particular file, statement within the file and view specification for initial display when arriving in the cited file.

A single, quickly executed command (Jump to Link) allows one to point at such a citation, or anywhere in the statement preceding the citation, and the system will go to the specific file and statement cited and show the associated material with the specified view parameters. This allows systems of interlinked documents and highly specific citations to be created.

A piece of the path through the chain of documents is saved so that one can return easily a limited distance back along his "trail," to previously referenced documents. Such a concept was originally suggested by Bush¹ in a fertile paper that has influenced our thinking in many ways.

Multiple windows

Another very useful feature is the ability to "split" the viewing screen horizontally and/or vertically in up to eight rectangular display windows of arbitrary size. Generally two to four windows are all that are used. Each window can contain a different view of the same or different locations, within the same or different files.³⁹

COLLABORATIVE DIALOGUE AND TELECONFERENCING

Introduction

The approach to collaboration support taken at ARC to date has two main thrusts:

- (1) Support for real-time dialogue (teleconferencing) for two or more people at two terminals who want to see and work on a common set of material. The collaborating parties may be further augmented with a voice telephone connection as well.
- (2) Support for written, recorded dialogue, distributed over time.

These two thrusts give a range of capabilities for support of dialogue distributed over time and space.

Teleconferencing support

Consider two people or groups of people who are geographically separated and who want to collaborate on a document, study a computer program, learn to use a new aspect of a system, or perform planning tasks, etc.

The workshop supports this type of collaboration by allowing them to link their terminals so that each sees the same information and either can control the system. This

function is available for both display and typewriter terminal users over the ARPANET.

The technique is particularly effective between displays because of the high speed of information output and the flexibility of being able to split the screen into several windows, allowing more than one document or view of a document to be displayed for discussion.

When a telephone link is also established for voice communication between the participants, the technique comes as close as any we know to eliminating the need for collaborating persons or small groups to be physically together for sophisticated interaction.

A number of other healthy approaches to teleconferencing are being explored elsewhere.^{11,12,16,17} It would be interesting to interface to such systems to gain experience in their use within workshops such as described here.

RECORDED DIALOGUE SUPPORT

Introduction

As ARC has become more and more involved in the augmentation of teams, serious consideration has been given to improving intra- and inter-team communication with whatever mixture of tools, conventions, and procedures will help.^{27,36,39}

If a team is solving a problem that extends over a considerable time, the members will begin to need help in remembering some of the important communications—i.e., some recording and recalling processes must be invoked, and these processes become candidates for augmentation.

If the complexity of the team's problem relative to human working capacity requires partitioning of the problem into many parts—where each part is independently attacked, but where there is considerable interdependence among the parts—the communication between various people may well be too complex for their own accurate recall and coordination without special aids.

Collaborating teams at ARC have been augmented by development of a "Dialogue Support System (DSS)," containing current and thoroughly used working records of the group's plans, designs, notes, etc. The central feature of this system is the ARC Journal, a specially managed and serviced repository for files and messages.

The DSS involves a number of techniques for use by distributed parties to collaborate effectively both using general functions in the workshop and special functions briefly described below and more fully in Reference 39. Further aspects are described in the section on Workshop Utility Service.

Document or message submission

The user can submit an NLS file, a part of a file, a file prepared on another system in the ARPANET (document), or text typed at submission time (message)

to the Journal system. When submitted, a copy of the document or message is transferred to a read-only file whose permanent safekeeping is guaranteed by the Journal system. It is assigned a unique catalog number, and automatically cataloged. Later, catalog indices based on number, author, and "titleword out of context" are created by another computer process.

Nonrecorded dialogue for quick messages or material not likely to be referenced in the future is also permitted.

One can obtain catalog numbers ahead of time to interlink document citations for related documents that are being prepared simultaneously. Issuing and controlling of catalog numbers is performed by a Number System (an automatic, crash-protected computer process).

At the time of submission, the user can contribute such information as: title, distribution list, comments, keywords, catalog numbers of documents this new one supersedes (updates), and other information.

The distribution is specified as a list of unique identification terms (abbreviated) for individuals or groups. The latter option allows users to establish dialogue groups. The system automatically "expands" the group identification to generate the distribution list of the individuals and groups that are its members. Special indices of items belonging to subcollections (dialogue groups) can be prepared to aid their members in keeping track of their dialogue. An extension of the mechanisms available for group distribution could give a capability similar to one described by Turoff.¹⁷

Entry of identification information initially into the system, group expansion, querying to find a person's or group's identification, and other functions are performed by an Identification System.

Document distribution

Documents are distributed to a person in one, two, or all of three of the following ways depending on information kept by the Identification System.

- (1) In hardcopy through the U.S. or corporation mail to those not having online access or to those desiring this mode,
- (2) Online as citations (for documents) or actual text (for messages) in a special file assigned to each user.
- (3) Through the ARPANET for printing or online delivery at remote sites. This delivery is performed using a standard Network wide protocol.

Document distribution is automated, with online delivery performed by a background computer process that runs automatically at specified times. Printing and mailing are performed by operator and clerical support. With each such printed document, an address cover sheet is automatically printed, so that the associated printout pages only need to be folded in half, stapled, and stamped before being dropped in the mail.

Document access

An effort has been made to make convenient both online and offline access to Journal documents. The master catalog number is the key to accessing documents. Several strategically placed hardcopy master and access collections (libraries) are maintained, containing all Journal documents.

Automatic catalog-generation processes generate author, number, and titleword indices, both online and in hardcopy.³⁸ The online versions of the indices can be searched conveniently with standard NLS retrieval capabilities.^{37,39,41}

Online access to the full text of a document is accomplished by using the catalog number as a file name and loading the file or moving to it by pointing at a citation and asking the system to "jump" there as described earlier.

SOFTWARE ENGINEERING AUGMENTATION SYSTEM

Introduction

One of the important application areas in ARC's work is software engineering. The economics of large computer systems, such as NLS, indicate that software development and maintenance costs exceed hardware costs, and that software costs are rising while hardware costs are rapidly decreasing. The expected lifetime of most large software systems exceeds that of any piece of computer hardware. Large software systems are becoming increasingly complex, difficult to continue evolving and maintain. Costs of additional enhancements made after initial implementation generally exceed the initial cost over the lifetime of the system. It is for these reasons that it is important to develop a powerful application area to aid software engineering. Areas of software engineering in which the ARC workshop offers aids are described below.

Design and review collaboration

During design and review, the document creation, editing, and studying capabilities are used as well as the collaboration, described above.

Use of higher level system programming languages

Programming of NLS is performed in a higher level ALGOL-like system programming language called L-10 developed at ARC. The L-10 language compiler takes its input directly from standard NLS structured files. The PDP-10 assembler also can obtain input from NLS files.

It is planned to extend this capability to other languages, for example, by providing an interface to the BASIC system available in our machine for knowledge workers wishing to perform more complex numerical tasks.

We are involved with developing a modular runtime-linkable programming system (MPS), and with planning a redesign of NLS to utilize MPS capabilities, both in cooperation with the Xerox Palo Alto Research Center. MPS will:

- (1) Allow a workshop system organization that will make it easier for many people to work on and develop parts of the same complex system semi-independently.
- (2) Make it easier to allow pieces of the system to exist on several processors.
- (3) Allow individual users or groups of users to tailor versions of the system to their special needs.
- (4) Make it easier to move NLS to other computers since MPS is written in itself.
- (5) Speed system development because of MPS's improved system building language facilities, integrated source-level debugging, measurement facilities, the ability to construct new modules by combining old ones, and to easily modify the system by changing module interconnection.

System documentation and source-code creation

Source-code creation uses the standard NLS hierarchical file structures and allows documentation and other programming conventions to be established that simplify studying of source-code files.

Debugging

A form of source-level debugging is allowed through development of several tools, of which the following are key examples:

- (1) A user program compilation and link loading facility that allows new or replacement programs to be linked into the running system to create revised versions for testing or other purposes.
- (2) NLS-DDT, a DDT like debugging facility with a command language more consistent with the rest of NLS, and simplifies display of system variables and data structures, and allows replacement of system procedures by user supplied procedures.
- (3) Use of several display windows so as to allow source code in some windows and control of DDT in others for the setting of breakpoints and display of variables and data structures.

Measurement and analysis

A range of measurement tools has been developed for analyzing system operation. These include the following:

- (1) Capabilities for gathering and reporting statistics on many operating system parameters such as utili-

zation of system components in various modes, queue lengths, memory utilization, etc.

- (2) The ability to sample the program counter for intervals of a selectable area of the operating system or any particular user subsystem to measure time spent in the sampled areas;
- (3) Trace and timing facilities to follow all procedure calls during execution of a specified function.
- (4) The ability to study page-faulting characteristics of a subsystem to check on its memory use characteristics.
- (5) The ability to gather NLS command usage and timing information.
- (6) The ability to study user interaction on a task basis from the point of view of the operating-system scheduler.
- (7) The ability to collect sample user sessions for later playback to the system for simulated load, or for analysis.

Maintenance

Maintenance programmers use the various functions mentioned above. The Journal is used for reporting bugs; NLS structured source code files simplify the study of problem areas and the debugging tools permit easy modification and testing of the modifications.

THE ARPA NETWORK INFORMATION CENTER (NIC)

Introduction

The NIC is presently a project embedded within ARC.³⁹ Workshop support for the NIC is based on the capabilities within the total ARC workshop system.

As useful as is the bootstrapping strategy mentioned earlier, there are limits to the type of feedback it can yield with only ARC as the user population. The NIC is the first of what we expect will be many activities set up to offer services to outside users. The goal is to provide a useful service and to obtain feedback on the needs of a wider class of knowledge workers. Exercised within the NIC are also prototypes of information services expected to be normal parts of the workshop.

The NIC is more than a classical information center, as that term has come to be used, in that it provides a wider range of services than just bibliographic and "library" type services.

The NIC is an experiment in setting up and running a general purpose information service for the ARPANET community with both online and offline services. The services offered and under development by the NIC have as their initial basic objectives:

- (1) To help people with problems find the resources (people, systems, and information) available within the network community that meet their needs.

- (2) To help members of geographically distributed groups collaborate with each other.

Following are the NIC services now provided to meet the above goals in serving the present clientele:

Current online services

- (1) Access to the typewriter version (TNLS) and display version (DNLS) of the Augmentation Research Center's Online System (NLS) for communicate creation, access, and linking between users, and for experimental use for any other information storage and manipulation purpose suitable for NLS and useful to Network participants.
- (2) Access to Journal, Number, and Identification Systems to allow messages and documents to be transmitted between network participants.
- (3) Access to a number of online information bases through a special Locator file using NLS link mechanisms and through a novice-oriented query system.

Current offline services

- (1) A Network Information Center Station set up at each network site.
- (2) Techniques for gathering, producing and maintaining data bases such as bibliographic catalogs, directories of network participants, resource information, and user guides.
- (3) Support of Network dialogue existing in hardcopy through duplication, distribution, and cataloging.
- (4) General Network referral and handling of document requests.
- (5) Building of a collection of documents potentially valuable to the Network Community. Initial concentration has been on obtaining documents of possible value to the Network builders.
- (6) As yet primitive selective document distribution to Station Collections.
- (7) Training in use of NIC services and facilities.

Conclusion

The Network Information Center is an example prototype of a new type of information service that has significant future potential. Even though it is presently in an experimental and developmental phase, it is providing useful online and offline services to the ARPANET community.

PLANS FOR A WORKSHOP UTILITY SERVICE

Motivation

It is now time for a next stage of application to be established. We want to involve a wider group of people

so that we can begin to transfer the fruits of our past work to them and with their assistance, to others, and so that we can obtain feedback needed for further evolution from wider application than is possible in our project alone.²⁸ We want to find and support selected groups who are willing to take extra trouble to be exploratory, but who:

- (1) Are not necessarily oriented to being core-workshop developers (they have their own work to do).
- (2) Can see enough benefit from the system to be tried and from the experience of trying it so that they can justify the extra risk and expense of being "early birds."
- (3) Can accept assurance that system reliability and stability, and technical/application help will be available to meet their conditions for risk and cost.

ARC is establishing a Workshop Utility Service, and promoting the type of workshop service described above as part of its long-term commitment to pursue the continued development of augmented knowledge workshops in a pragmatic, evolutionary manner.

It is important to note that the last few years of work have concentrated on the means for delivering support to a distributed community, for providing teleconferencing and other basic processes of collaborative dialogue, etc. ARC has aimed consciously toward developing experience and capabilities especially applicable to support remote and distributed groups of exploratory users for this next stage of wider-application bootstrapping.

One aspect of the service is that it will be an experiment in harnessing the new environment of a modern computer network to increase the feasibility of a wider community of participants cooperating in the evolution of an application system.

Characteristics of the planned service

The planned service offered will include:

- (1) Availability of Workshop Utility computer service to the user community from a PDP-10 TENEX system operated by a commercial supplier.
- (2) Providing training as appropriate in the use of Display NLS (DNLS), Typewriter NLS (TNLS), and Deferred Execution (DEX) software subsystems.
- (3) Providing technical assistance to a user organization "workshop architect" in the formulation, development, and implementation of augmented knowledge work procedures within selected offices at the user organization.⁶

This assistance will include help in the development of NLS use strategies suitable to the user environments, procedures within the user organization for implementing these strategies, and possible special-application NLS extensions (or

simplifications) to handle the mechanics of particular user needs and methodologies.

- (4) Providing "workshop architect" assistance to help set up and assist selected geographically distributed user groups who share a special discipline or mission orientation to utilize the workshop utility services and to develop procedures, documentation, and methodology for their purposes.

GENERAL DESCRIPTION OF SOME WORKSHOP UTILITY SUBSYSTEMS

Introduction

Within a particular professional task area (mission- or discipline-oriented) there are often groups who could be benefited by using special workshop subsystems. These subsystems may be specialized for their specific application or research domain or for support of their more general knowledge work. Our goal is to offer a workshop utility service that contains a range of subsystems and associated methodology particularly aimed at aiding general knowledge work, and that also supports in a coordinated way special application subsystems either by interfacing to subsystems already existing, or by developing new subsystems in selected areas.

In the descriptions to follow are a number of workshop subsystem domains that are fundamental to a wide range of knowledge work in which ARC already has extensive developments or is committed to work. For each subsystem we include some general comments as well as a brief statement of current ARC capabilities in the area.

Document development, production, and control

Here a system is considered involving authors, editors, supervisors, typists, distribution-control personnel, and technical specialists. Their job is to develop documents, through successive drafts, reviews, and revisions. Control is needed along the way of bibliography, who has checked what point, etc. Final drafts need checkoff, then production. Finally distribution needs some sort of control. If it is what we call a "functional document" such as a user guide, then it needs to be kept up to date.³⁹ There is a further responsibility to keep track of who needs the documents, who has what version, etc.

Within the ARC workshop, documents ranging from initial drafts to final high-quality printed publications can be quickly produced with a rich set of creation and editing functions. All of ARC's proposals, reports, designs, letters, thinkpieces, user documentation, and other such information are composed and produced using the workshop.

Documents in a proof or finished form can be produced with a limited character set and control on a line printer or typewriter, or publication-quality documents can be produced on a photocomposer microfilm unit.

Presently there are on the order of two hundred special directives that can be inserted in text to control printing. These directives control such features as typefont, pagination, margins, headers, footers, statement spacing, typefont size and spacing, indenting, numbering of various hierarchical levels, and many other parameters useful for publication quality work. Methodology to perform the creation, production, and controlling functions described above has been developed, although much work at this level is still needed.

In terms of future goals, one would like to have display terminals with a capability for the range of fonts available on the photocomposer so that one could study page layout and design interactively, showing the font to be used, margins, justification, columnization, etc. on the screen rather than having to rely on hardcopy proof-sheets.

To prepare for such a capability, plans are being made to move toward an integrated portrayal mechanism for both online and hardcopy viewing.

Collaborative dialogue and teleconferencing

Effective capabilities have already been developed and are in application, as discussed above. There is much yet to do. The Dialogue Support System will grow to provide the following additional general online aids:

Link-setup automation; back-link annunciators and jumping; aids for the formation, manipulation, and study of sets of arbitrary passages from among the dialogue entries; and integration of cross-reference information into hardcopy printouts. Interfaces will probably be made to other teleconferencing capabilities that come into existence on the ARPANET.

It also will include people-system developments: conventions and working procedures for using these aids effectively in conducting collaborative dialogue among various kinds of people, at various kinds of terminals, and under various conditions; working methodology for teams doing planning, design, implementation coordination; and so on.

Meetings and conferences

Assemblies of people are not likely for a long time, if ever, to be supplanted in total by technological aids. Online conferences are held at ARC for local group meetings and for meetings where some of the participants are located across the country.

Use is made of a large-screen projection TV system to provide a display image that many people in a conference room can easily see. This is controlled locally or remotely by participants in the meeting, giving access to the entire recorded dialogue data base as needed during the meeting and also providing the capability of recording real-time

meeting notes and other data. The technique also allows mixing of other video signals.

Management and organization

The capabilities offered in the workshop described in this paper are used in project management and administration.³⁹ Numerical calculations can also be performed for budget and other purposes, obtaining operands and returning results to NLS files for further manipulation.

Where an organization has conventional project management operations, their workshop can include computer aids for techniques such as PERT and CPM. We want to support the interfacing that our Core Workshop can provide to special application systems for management processes.

We are especially interested at this stage, in management of project teams—particularly, of application-systems development teams.

Handbook development

Capabilities described above are being extended toward the coordinated handling of a very large and complex body of documentation and its associated external references. The goal is that a project or discipline of ever-increasing size and complexity can be provided with a service that enables the users to keep a single, coordinated “superdocument” in their computer; that keeps up to date and records the state of their affairs; and provides a description of the state of the art in their special area.

Example contents would be glossaries, basic concept structure, special analytic techniques, design principles, actual design, and implementation records of all developments.

Research intelligence

The provisions within the Dialogue Support System for cataloging and indexing internally generated items also support the management for externally generated items, bibliographies, contact reports, clippings, notes, etc. Here the goal is to give a human organization (distributed or local) an ever greater capability for integrating the many input data concerning its external environment; processing (filtering, transforming, integrating, etc.) the data so that it can be handled on a par with internally generated information in the organization’s establishing of plans and goals; and adapting to external opportunities or dangers.³⁸

Computer-based instruction

This is an important area to facilitate increasing the skills of knowledge workers. ARC has as yet performed little direct work in this area. We hope in the future to work closely with those in the computer-based instruction

area to apply their techniques and systems in the workshop domain.

In training new and developing users in the use of the system, we have begun using the system itself as a teaching environment. This is done locally and with remote users over the ARPANET.

Software engineering augmentation

A major special application area described above, that has had considerable effort devoted to it, is support of software engineers. The software-based tools of the workshop are designed and built using the tools previously constructed. It has long been felt^{24,29} that the greatest “bootstrapping” leverage would be obtained by intensively developing the augmented workshop for software engineers, and we hope to stimulate and support more activity in this area.

Knowledge workshop analysis

Systematic analysis has begun of the workshop environment at internal system levels, at user usage levels, and at information-handling procedure and methodology levels. The development of new analytic methodology and tools is a part of this process. The analysis of application systems, and especially of core-workshop systems, is a very important capability to be developed. To provide a special workshop subsystem that augments this sort of analytic work is a natural strategic goal.

CONCLUSION—THE NEED FOR LONG-TERM COMMITMENT

As work progresses day-to-day toward the long-term goal of helping to make the truly augmented knowledge workshop, and as communities of workshop users become a reality, we at ARC frequently reflect on the magnitude of the endeavor and its long-term nature.²²

Progress is made in steps, with hundreds of short-term tasks directed to strategically selected subgoals, together forming a vector toward our higher-level goals.

To continue on the vector has required a strong commitment to the longer-range goals by the staff of ARC.

In addition, we see that many of the people and organizations we hope to enlist in cooperative efforts will need a similar commitment if they are to effectively aid the process.

One of ARC’s tasks is to make the long-term objectives of the workshop’s evolutionary development, the potential value of such a system, and the strategy for realizing that value clear enough to the collaborators we seek, so that they will have a strong commitment to invest resources with understanding and patience.

One key for meeting this need will be to involve them in serious use of the workshop as it develops. The plans for the Workshop Utility are partly motivated by this objective.

Although the present ARC workshop is far from complete, it does have core capabilities that we feel will greatly aid the next communities of users in their perception of the value of the improved workshops of the future.

ACKNOWLEDGMENTS

During the 10 year life of ARC many people have contributed to the development of the workshop described here. There are presently some 35 people—clerical, hardware, software, information specialists, operations researchers, writers, and others—all contributing significantly toward the goals described here.

The work reported here is currently supported primarily by the Advanced Research Projects Agency of the Department of Defense, and also by the Rome Air Development Center of the Air Force and by the Office of Naval Research.

REFERENCES

- Bush, V., "As We May Think," *Atlantic Monthly*, pp. 101-108, July 1945 (SRI-ARC Catalog Item 3973).
- Licklider, J. C. R., "Man-Computer Symbiosis," *IEEE Transactions on Human Factors in Electronics*, Vol. HFE-1, pp. 4-11, March, 1960 (SRI-ARC Catalog Item 6342).
- Drucker, P. F., *The Effective Executive*, Harper and Row, New York, 1967 (SRI-ARC Catalog Item 3074).
- Drucker, P. F., *The Age of Discontinuity—Guidelines to our Changing Society*, Harper and Row, New York, 1968 (SRI-ARC Catalog Item 4247).
- Dalkey, N., *The Delphi Method—An Experimental Study of Group Opinion*, Rand Corporation Memorandum RM-5888-PR, 1969 (SRI-ARC Catalog Item 3896).
- Allen, T. J., Piepmeier, J. M., Cooney, S., "Technology Transfer to Developing Countries—The International Technological Gatekeeper," *Proceedings of the ASIS*, Vol. 7, pp. 205-210, 1970 (SRI-ARC Catalog Item 13959).
- Roberts, L. G., Wessler, B. D., "Computer Network Development to Achieve Resource Sharing," *AFIPS Proceedings*, Spring Joint Computer Conference, Vol. 36, pp. 543-549, 1970 (SRI-ARC Catalog Item 4564).
- Roberts, L. G., Wessler, B. D., *The ARPA Network*, Advanced Research Projects Agency, Information Processing Techniques, Washington, D.C. May 1971 (SRI-ARC Catalog Item 7750).
- Bobrow, D. G., Burchfiel, J. D., Murphy, D. L., Tomlinson, R. S., "TENEX—A Paged Time Sharing System for the PDP-10," presented at *ACM Symposium on Operating Systems Principles*, October 18-20, 1971. Bolt Beranek and Newman Inc., August 15, 1971 (SRI-ARC Catalog Item 7736).
- Weinberg, G. M., *The Psychology of Computer Programming*, Van Nostrand Reinhold Company, New York, 1971 (SRI-ARC Catalog Item 9036).
- Hall, T. W., "Implementation of an Interactive Conference System," *AFIPS Proceedings*, Spring Joint Computer Conference, Vol. 38, pp. 217-229, 1971 (SRI-ARC Catalog Item 13962).
- Turoff, M., "Delphi and its Potential Impact on Information Systems," *AFIPS Proceedings*, Fall Joint Computer Conference, Vol. 39, pp. 317-326, 1971 (SRI-ARC Catalog Item 7966).
- Roberts, L. G., *Extensions of Packet Communication Technology to a Hand Held Personal Terminal*, Advanced Research Projects Agency, Information Processing Techniques, January 24, 1972 (SRI-ARC Catalog Item 9120).
- Kahn, R. E., "Resource-Sharing Computer Communication Networks," *Proceedings of the IEEE*, Vol. 147, pp. 147, September 1972 (SRI-ARC Catalog Item 13958).
- Anderson, R. E., Coover, E. R., "Wrapping Up the Package—Critical Thoughts on Applications Software for Social Data Analysis," *Computers and Humanities*, Vol. 7, No. 2, pp. 81-95, November 1972 (SRI-ARC Catalog Item 13956).
- Lipinski, A. J., Lipinski, H. M., Randolph, R. H., "Computer Assisted Expert Interrogation—A Report on Current Methods Development," *Proceedings of First International Conference on Computer Communication*, Winkler, Stanley (ed), October 24-26, 1972, Washington, D.C., pp. 147-154 (SRI-ARC Catalog Item 11980).
- Turoff, M., "'Party-Line' and 'Discussion' Computerized Conference Systems," *Proceedings of First International Conference on Computer Communication*, Winkler, Stanley (ed), October 24-26, 1972, Washington, D. C., pp. 161-171, (SRI-ARC Catalog Item 11983).
- Licklider, J. C. R., Taylor, R. W., Herbert, E., "The Computer as a Communication Device," *International Science and Technology*, No. 76, pp. 21-31, April 1968 (SRI-ARC Catalog Item 3888).
- Engelbart, D. C., "Augmenting your Intellect," (Interview with D. C. Engelbart), *Research/Development*, pp. 22-27, August 1968 (SRI-ARC Catalog Item 9698).
- Haavind, R., "Man-Computer 'Partnerships' Explored," *Electronic Design*, Vol. 17, No. 3, pp. 25-32, February 1, 1969 (SRI-ARC Catalog Item 13961).
- Field, R. K., "Here Comes the Tuned-In, Wired-Up, Plugged-In, Hyperarticulate Speed-of-Light Society—An Electronics Special Report: No More Pencils, No More Books—Write and Read Electronically," *Electronics*, pp. 73-104, November 24, 1969 (SRI-ARC Catalog Item 9705).
- Lindgren, N., "Toward the Decentralized Intellectual Workshop," *Innovation*, No. 24, pp. 50-60, September 1971 (SRI-ARC Catalog Item 10480).
- Engelbart, D. C., "Special Considerations of the Individual as a User, Generator, and Retriever of Information," *American Documentation*, Vol. 12, No. 2, pp. 121-125, April 1961 (SRI-ARC Catalog Item 585).
- Engelbart, D. C., "A Conceptual Framework for the Augmentation of Man's Intellect," *Vistas in Information Handling*, Howerton and Weeks (eds), Spartan Books, Washington, D.C., 1963, pp. 1-29 (SRI-ARC Catalog Item 9375).
- English, W. K., Engelbart, D. C., Berman, M. A., "Display-Selection Techniques for Text Manipulation," *IEEE Transactions on Human Factors in Electronics*, Vol. HFE-8, No. 1, pp. 5-15, March 1967 (SRI-ARC Catalog Item 9694).
- Engelbart, D. C., English, W. K., "A Research Center for Augmenting Human Intellect," *AFIPS Proceedings*, Fall Joint Computer Conference, Vol. 33, pp. 395-410, 1968 (SRI-ARC Catalog Item 3954).
- Engelbart, D. C., "Intellectual Implications of Multi-Access Computer Networks," paper presented at *Interdisciplinary Conference on Multi-Access Computer Networks*, Austin, Texas, April 1970, preprint (SRI-ARC Journal File 5255).

BY OTHER PEOPLE, WITH SUBSTANTIVE DESCRIPTION OF ARC DEVELOPMENTS

OPEN-LITERATURE ITEMS BY ARC STAFF

28. Engelbart, D. C., *Coordinated Information Services for a Discipline- or Mission-Oriented Community*, Stanford Research Institute Augmentation Research Center, December 12, 1972 (SRI-ARC Journal File 12445). Also published in "Time Sharing—Past, Present and Future," *Proceedings of the Second Annual Computer Communications Conference* at California State University, San Jose, California, January 24-25, 1973, pp. 2.1-2.4, 1973. Catalog Item 5139).

RELEVANT ARC REPORTS

29. Engelbart, D. C., *Augmenting Human Intellect—A Conceptual Framework*, Stanford Research Institute Augmentation Research Center, AFOSR-3223, AD-289 565, October 1962 (SRI-ARC Catalog Item 3906).
30. Engelbart, D. C., Huddart, B., *Research on Computer-Augmented Information Management (Final Report)*, Stanford Research Institute Augmentation Research Center, ESD-TDR-65-168, AD 622 520, March 1965 (SRI-ARC Catalog Item 9690).
31. Engelbart, D. C., *Augmenting Human Intellect—Experiments, Concepts, and Possibilities—Summary Report*, Stanford Research Institute Augmentation Research Center, March 1965 (SRI-ARC Catalog Item 9691).
32. English, W. K., Engelbart, D. C., Huddart, B., *Computer Aided Display Control—Final Report*, Stanford Research Institute Augmentation Research Center, July 1965 (SRI-ARC Catalog Item 9692).
33. Engelbart, D. C., English, W. K., Rulifson, J. F., *Development of a Multidisplay*, Time-Shared Computer Facility and Computer-Augmented Management-System Research, Stanford Research Institute Augmentation Research Center, AD 843 577, April 1968 (SRI-ARC Catalog Item 9697).
34. Engelbart, D. C., *Human Intellect Augmentation Techniques—Final Report*, Stanford Research Institute Augmentation Research

Center, CR-1270 N69-16140, July 1968 (SRI-ARC Catalog Item 3562).

35. Engelbart, D. C., English, W. K., Evans, D. C., *Study for the Development of Computer Augmented Management Techniques—Interim Technical Report*, Stanford Research Institute Augmentation Research Center, RADC-TR-69-98, AD 855 579, March 8, 1969 (SRI-ARC Catalog Item 9703).
36. Engelbart, D. C., SRI-ARC Staff, *Computer-Augmented Management-System Research and Development of Augmentation Facility—Final Report*, Stanford Research Institute Augmentation Research Center, RADC-TR-70-82, April 1970 (SRI-ARC Catalog Item 5139).
38. Engelbart, D. C., *Experimental Development of a Small Computer-Augmented Information System—Annual Report*, Stanford Research Institute Augmentation Research Center, April 1972 (SRI-ARC Catalog Item 10045).
39. *Online Team Environment—Network Information Center and Computer Augmented Team Interaction*, Stanford Research Institute Augmentation Research Center, RADC-TR-72-232, June 8, 1972 (SRI-ARC Journal File 13041).

RELEVANT ARTICLES IN ARC/NIC JOURNAL

40. Engelbart, D. C., *SRI-ARC Summary for IPT Contractor Meeting*, San Diego, January 8-10, 1973, Stanford Research Institute Augmentation Research Center, January 7, 1973 (SRI-ARC Journal File 13537).

MOVIE AVAILABLE FROM ARC FOR LOAN

41. *Augmentation of the Human Intellect—A Film of the SRI-ARC Presentation at the 1969 ASIS Conference*, San Francisco (A 3-Reel Movie, Stanford Research Institute Augmentation Research Center, October 1969 (SRI-ARC Catalog Item 9733).