
Building an Electronic Community System

BRUCE R. SCHATZ

BRUCE R. SCHATZ is Director of the Community Systems Laboratory and is a research scientist in molecular and cellular biology and in MIS at the University of Arizona. He spent ten years in industrial research laboratories at Bell Labs and Bellcore as a lead architect on a variety of research and development projects in information and communications systems. He came to the University of Arizona in 1989 to found a laboratory to construct novel science information systems and propagate them to large-scale communities. The funding for this laboratory includes a grant (on which he was Principal Investigator) that was one of the winners of the NSF National Collaboratory competition. He earned an M.S. in artificial intelligence from Massachusetts Institute of Technology and a Ph.D. in computer science from the University of Arizona. His research interests include community systems, electronic libraries, computational biology, and network applications.

ABSTRACT: An *electronic community system* encodes and manipulates the range of knowledge and values necessary to function effectively in a community or organization. The knowledge includes both formal data and literature and informal results and news. The manipulation includes both browsing through the available knowledge, and recording and sharing interrelationships between the items. A large-scale experiment is underway to build an electronic community system for the community of scientists studying the nematode worm *C. elegans*, a model organism in molecular biology. This paper discusses a model for community systems and previous such systems in science, the biology experiment and a previous system, the enabling technology for handling the knowledge, the enabling mechanisms for handling the values, the state of the prototype, and speculations on future applications in supporting organizational memory.

KEY WORDS AND PHRASES: electronic community systems, electronic communities, scientific applications, information spaces, telesophy, organizational memory.

TO FUNCTION EFFECTIVELY IN AN ORGANIZATION, one needs access to a wide variety of knowledge. This knowledge includes not only the archives of financial data about

An earlier version of this paper was originally published in the *Proceedings of the Twenty-Fourth Hawaii International Conference on System Sciences* (IEEE Computer Society Press, 1991). This paper is partially based upon a grant proposal written with Samuel Ward, who introduced me to the worm community and has supported me at the University of Arizona. Many people worked on the implementation of the prototype, including Terry Friedman, Ed Grossman, Scott Hudson, Kevin Powell, and Andrey Yeatts. Mark Edgley and John Sulston provided much of the worm data. John Wooley provided helpful advice. Planning for this project was supported by a Small Grant for Exploratory Research DIR-9003540 from the National Science Foundation. The project itself is now supported by NSF IRI-9015407 through the Special Initiative in Coordination Theory and Collaboration Technology.

Journal of Management Information Systems / Winter 1991-92, Vol. 8, No. 3, pp. 87-107.

Copyright © M.E. Sharpe, Inc., 1992.

products and technical information about designs, but also test results, meeting reports, and other informal sources. To integrate this knowledge, one must understand the relationships between the available items in the context of the company values and culture. Productivity would be greatly enhanced if a substantial fraction of this integrated knowledge were easily available from the employee's desktop computer. Traditional management information systems have only supported a small portion of this functionality.

It is now possible to build substantial prototypes of information systems that handle such a wide variety of integrated knowledge. These can be termed *electronic community systems*, which encode and manipulate formal and informal knowledge and their interrelationships. A large-scale experiment is being carried out to encode the knowledge of a community of scientists and to build a software environment to manipulate this knowledge from their laboratories. The community includes those scientists studying the nematode worm *C. elegans*, a model organism in molecular biology. This paper discusses a model for community systems and previous such systems in science, the biology experiment and a previous system, the enabling technology for handling the knowledge, the enabling mechanisms for handling the values, the state of the prototype, and speculations on future applications in supporting organizational memory.

A Model for Electronic Community Systems

THE WORD "COMMUNITY" IS CLOSELY ALLIED WITH THE WORDS "common," meaning "the same," and "communicate," meaning "to exchange information." Originally, *community* referred to the people residing in some small physical location and more generally to their shared values. The meaning of the word has been extended to groups of people with common interests and shared values who may reside in geographically separate places. Thus, one may refer to the scientific community or the physics community or the relativity community. This section discusses a general model for a community and its support by a computer system.

To support a community electronically, it is necessary to encode as much of its knowledge as possible. Figure 1 illustrates the range of possible knowledge that might be supported by a computer system. To live effectively within a community, one must have available both the formal archival material and the informal transient folklore. This includes the fundamental items of data for the community, for example, as maintained by database management systems, and the intermediate results, for example, as contained in electronic mail messages. This includes the archival literature for the community, for example, as maintained by information retrieval systems, and the intermediate news, for example, as contained in electronic bulletin boards. Finally, it includes support for the shared values as well as the common interests. The mores of the community can be supported by means of a variety of mechanisms for recording the relationships between the data and information, for example, by providing hypertext documents.

An *electronic community system* is a computer system that encodes the knowledge of a community and provides an environment that supports manipulation of that




Figure 1. An Electronic Community System Records and Manipulates All the Knowledge of a Community or Organization

knowledge. Different communities have different knowledge but their environments have great similarities. The community knowledge might be thought of as being stored in an electronic library. Much of the material originates within external sources. The environment must accordingly provide software for building a library to access these sources, for example, convenient mechanisms for encoding and browsing what is available. But, unlike existing physical libraries, a community library is dynamic and the members will actively add items to it. The environment must also provide software for updating this library, that is, convenient mechanisms for refereeing and sharing added items. The environment thus provides support for both the knowledge and the mores of the community.

The functionality of an electronic community system can be motivated by considering the use of a physical library. Consider the analogy of doing research in a physical library in order to write a book. You start with references from a paper or colleagues, look the references up in the card catalog, go to an appropriate section of the library, and scan your eye along the titles on the spines of the books. If any books look relevant, you pull them off the shelf for detailed examination of the pages. If some pages look relevant, you make a copy for later use. After some scanning and examining, you go to another section of the library, often using references found in the previous section. When the research phase is finished, you write your book, utilizing references to copied pages, and submit your book to be published (and subsequently placed itself into the library).

In the model of a community library, the books are distributed multimedia objects. There are three basic stages in the interaction process: browsing, filtering, and sharing. In *browsing*, the user can rapidly examine the items in the library. This can be accomplished via search, by giving an associative specification and viewing matching items, or via navigation, by following the connections from a given item. The results can be displayed at a variety of summary levels. Multiple searches and navigations can be issued and cross-compared to located desired items.

The next stage is *filtering*, where the user culls the items located by browsing into

some desired set, relevant to the current need. If the browsing speed is sufficiently fast, user view of displays may be sufficient to select relevant items manually. If the items are too numerous or too complex, manual examination may not be sufficient. In this case, a set of selected items can be passed into an external analysis program for automatic filtering. Such a program might sort the items by date or perform a complicated computation to determine rank ordering against some similarity metric.

After a set of desired existing items have been found, the user may wish to add this set with comments back into the library. *Sharing* is the support for publishing in the electronic library. A variety of mechanisms are supported within an electronic community system for grouping the items to record their relationships, for example, storing a set with a description of its relationship or forging a connection link between related items. Mechanisms are also supported for writing hyperdocuments, which incorporate other items into the text via embedded links. Once a sufficiently important new group or document has been composed, facilities are available for releasing this to the community. A variety of mechanisms are supported to provide editorial and privacy control of the release process. These mechanisms are the attempt to encode the mores of the community, by permitting members of the community to control the quality of the material in the library and who may view the material.

Electronic Scientific Communities

THE BUSINESS OF SCIENTIFIC RESEARCH IS AN UNUSUALLY GOOD DOMAIN to investigate the development of an electronic community system. Practicing scientists need access to a wide variety of knowledge to carry out their research. Much of this knowledge resides in formal published literature, but much also resides in informal community knowledge. Some of this informal community knowledge, such as preliminary results, will eventually become published, but other knowledge, such as details of methods and the "lore" of experimental systems, never reaches the formal literature. The scientist who shares a community's current informal knowledge and has rapid access to the formal knowledge can do better research. This is particularly true in the biological sciences, which are largely data-driven, because the choice and design of experiments depend on familiarity with the most current methods and knowledge. As the pace of science increases, only a small number of insiders who lead each field, the "invisible college" [2], have enough knowledge to perform seminal research efficiently. If the informal community knowledge could be captured and disseminated more widely, the quality and efficiency of scientific research would improve. This is because the invisible college would be larger and because it would be open to scientists from diverse disciplines, which would encourage novel interdisciplinary research.

The existence of nationwide networks has fostered electronic scientific communities. The ARPANET was the first nationwide computer network widely available to the scientific community in the United States. It was constructed during the late 1960s and reached its potential throughout the 1970s. The original motivation for the network was to support remote access to the large "supercomputers" constructed and purchased

by ARPA-funded researchers. What emerged as the most important service, however, was the new facility of *electronic mail*, which provided a new communications medium. The ability to convey informal information rapidly caused a new feeling of closeness among the researchers on the network and the emergence of the first widespread electronic community, the ARPANET community [7]. Researchers on the ARPANET could get essential information more quickly than those not on the net and many collaborations took place without the collaborators ever physically meeting.

Early users of electronic mail in the ARPANET noted the convenience of being able to send items to many others on a distribution list. Standardized lists became established to distribute messages to people interested in a wide variety of topics. These lists evolved into the next generation of community system, the *electronic bulletin board*. An illustrative current-day bulletin board system is Netnews [10], which distributes messages over USENET [9]. USENET is not a centrally planned and maintained network, but a loose collection of computers running the Unix™ operating system connected by a wide variety of physical transmission lines from high-speed leased lines to ordinary telephone lines. It operates today with more than 250,000 users on more than 10,000 machines spread throughout the world. Netnews contains more than 650 boards across a wide variety of topics, ranging from comments about existing computers to technical science to popular culture to job positions to movie reviews to cooking recipes. The software functionality has evolved to support streamlined posting to the appropriate boards, comments on previous messages, reading of selected boards, and saving of selected messages.

Everyday operation of Netnews shows the benefits of community sharing. When you post a technical question, you often get a detailed technical answer from somewhere out in the Net, often from a place you would never have thought of looking. Frequently, your posting stimulates a series of postings, each illuminating the problem from a different point of view. It is common to see the understanding of a problem evolve over a week through comments from a series of different postings from different parts of the world. For example, a recent query by the author concerning radio interference on laptop computers elicited helpful comments and critiques from people at sites in Massachusetts, California, Wisconsin, Oregon, Hawaii, Ontario, Germany, and Sweden. There is a real feeling of community interaction to solve shared problems on the many responsible boards. As with the ARPANET boards, Netnews tends to be self-policing. The users tend to be responsible individuals who understand that the system is supported by the generosity of their employers. People who abuse the Net (by posting inflammatory or irresponsible messages) are quickly dealt with by peer pressure. There are elaborate documents on appropriate Net etiquette: what boards are suitable for what topics, what content is appropriate for a posting, how to be terse and polite, and so on. Different board types have evolved a spectrum of editorial control, ranging from boards where anyone can post anything to ones where all messages pass through a human moderator for topic and quality control.

As the speed of networks increases, so does the range of information that can be effectively encoded within an electronic community. The dream of an all-encompassing science information system is an old one, since the possibility of being able to sit

in front of a computer and be able to access all the knowledge you need for your research is so attractive. This dream has resurfaced periodically whenever the computing and communications technology makes a dramatic increase in functionality. See, for example, the "future of libraries" study after the advent of minicomputers in the 1960s [6], the "world scientific information system" study after the advent of computer networks in the 1970s [16], and the "national collaboratory" study after the advent of workstations and supercomputer networks in the 1980s [8]. The forthcoming NREN (National Research and Educational Network) will provide network speeds fast enough to support interactive manipulation of a wide variety of material across the national scientific network. This leads to the possibility of realizing an all-encompassing information system with the next generation of community systems.

Building an Electronic Community System

COMMUNITY SYSTEMS OF THE NEAR FUTURE will support the complete range of knowledge and functionality discussed in the Model section above. They will support a wide variety of database management and information retrieval functions to support a wide variety of formal experimental data and literature information. In addition, they will support a wide variety of electronic mail and bulletin board functions to support a wide variety of informal results and news.

The Community Systems Laboratory is building an electronic community system in the domain of scientific research and evaluating its use within the community as a large-scale experiment. The resulting system is meant to serve a wide variety of communication needs within the community, both retrieval and analysis, as well as rapid sharing of knowledge with others. It will permit researchers, who have common interests and shared values but are geographically dispersed, to browse and share the community knowledge. The scientific community chosen for this experimental project is the community of molecular biologists united by their common use of a model organism, the nematode worm *Caenorhabditis elegans* [11].

The Worm Community

Building an electronic scientific community in today's largely nonelectronic world requires a specialized community with an appropriate set of characteristics. It must have a large amount of data, both formal and informal, and a real need to manipulate these data extensively. The data must be freely available and already largely in electronic form. There must be many interested users who are willing to experiment with new technology and who have adequate computer equipment and network connections. There must be real support for data administration and software development, which implies that the community must be an important one scientifically so that adequate funding is available.

A scientific community that exhibits these properties is the *worm community*, the molecular biologists who utilize the nematode worm *Caenorhabditis elegans*. Molecular biology is a largely data-driven experimental science and, due to such efforts as

the Human Genome Initiative, its data are growing rapidly and being stored in databases. Communities in molecular biology often form around organisms, rather than techniques or problems. *C. elegans* is a nonparasitic worm found in the soil, which has been extensively studied, with a wide range of experimental data available on its genetics, anatomy, and development [17]. The "worm" has become a primary model organism and will likely become one of the first to be completely sequenced.

The worm community itself is young, but growing rapidly, with more than 500 researchers at the last large meeting in June 1991. It has a close-knit and communicative group of "insiders," the postdoctoral fellows of Sydney Brenner who initiated the modern genetics of *C. elegans* in the late 1960s. It has a strong community tradition of free sharing of unpublished data, unlike the competitive nature of many other areas of science. Substantial bodies of data are already in electronic form, and there is an adequate level of computer literacy and interest in electronic communication.

Two examples of widely shared unpublished data within the worm community illustrate its suitability as a cooperative experimental electronic community. The *Worm Breeder's Gazette* is a newsletter analogous to a moderated electronic bulletin board, which consists of short research items and has been published several times a year for more than ten years in an unrefereed open format. The physical map database is an electronic recording of an ordered set of cloned DNA fragments that constitute the genetic material of the worm [1]. Its curators map new fragments and distribute the database as a service to the community; the easy availability has dramatically facilitated molecular analysis of genes.

The worm community is an excellent candidate for an electronic scientific community due to a number of unusual characteristics. It is the right size—big enough so that newcomers no longer know the pioneers directly but small enough so that fierce competition has not yet set in. It is the right age—old enough so that there is extensive knowledge already discovered but young enough so that the insiders still remember the original days and want to preserve their closeness. It has the right importance—significant enough so that discoveries make a scientific difference but offbeat enough so that researchers do not hoard their data.

Finally, the worm community has always had a special tradition of sharing knowledge. Reasons include the fact that many members of the worm community were trained in the phage group where openness was encouraged and that there has always been a primary goal of understanding everything about the worm. The worm community's informal network of communication is becoming inadequate as the community grows, and there is concern among the insiders about losing the community's unique flavor. Thus, there is an immediate need for an electronic worm community and a favorable set of conditions for an experiment in electronic communication.

Worm Community Knowledge

One major advantage of the worm community as a community system testbed is the large amount of important data available. The most significant available sources are

listed below. The categories for these materials span the range of editorial control: published literature is refereed; archival data are edited (checked for quality); informal information is moderated (checked for topic); and unpublished data are posted (no checking). Archival data are typically maintained by a central curator, whereas unpublished data are maintained by individual researchers. Similarly, much of the informal information is unrefereed literature.

Archival Data

Gene Descriptions. These are text descriptions of the phenotypes of worm mutations and other genetic effects. About 1,200 of perhaps 5,000 genes are known.

Genetic Map. This represents the relative positions of genes on the chromosomes, based on how often two genes recombine during sexual reproduction. The display is a line drawing with marks denoting the gene positions.

Physical Map. This represents cloned fragments of worm DNA and how they overlap to form the chromosomes. Genes known to be within a fragment thus have a known absolute position. The display is a drawing with many overlapping lines and located genes. About 11,000 clones nearly cover the chromosomes.

DNA Sequences. These are the nucleotide codes for the genes, with only a few known at present. They display as strings of ACGT.

Cell List. This table and text identify every one of the 959 cells and their location in the worm.

Wiring Diagram. This table gives the complete connection pattern of all 302 neurons. It is displayed as a line drawing.

Cell Lineage. This table and line drawing give the complete developmental history of every cell, i.e., which cell develops into which other cells during growth.

Formal Literature

Bibliography. This text gives citations for most of the *C. elegans* literature. The current list of some 1,400 articles is maintained by a central curator.

Abstracts. Text abstracts for most of these articles are available in on-line literature databases such as Medline, Biosis, and Agricola.

Full-text. Text of the complete contents of all the papers would be preferable for a complete community library. If copyright permission were available, it would be possible to scan and recognize the characters in worm journal articles.

Page Images. The most desirable storage would be complete article contents, including drawings, figures, tables, graphs, and equations, in addition to text. The most practical approach is to digitally scan page images of the nontextual material. After an article is located by searching the full-text, the text can be displayed in formatted form along with the images of the figures.

Worm Book. This is the standard reference book *The Nematode Caenorhabditis Elegans* [17], which contains review articles on all aspects of *C. elegans* biology as well as the existing data, nearly 700 pages total.

Informal Literature

Newsletters. The *Worm Breeder's Gazette* has been published for ten years and is a rich source of one-page notes about items of interest from experimental results to methods to news.

Conferences. The biannual *C. elegans* Meeting publishes one-page abstracts of presented papers that can be cited as personal communications.

Unpublished Data

Lab Directory. This is the list of current researchers with their addresses.

Strain List. This is the list of the worm mutations available from the centralized stock center.

Useful unpublished data are also maintained in the laboratories of individual researchers. They include: experimental methods (text), genomic maps (drawings), micrographs (images), and strain lists (text).

Analysis Software

There are a variety of programs available for analyzing the biological data, such as comparing sequence similarity. The environment provides a facility for selecting a set of items and passing this set into an external analysis program. Some of these programs also provide sophisticated displays of the data, for example, genetic and physical maps.

Community Lore

Much of the knowledge about the worm is not currently recorded anywhere. The Worm Community System will support facilities for entering annotations, specifying relationship links, and writing documents. This new material will be added to the shared community library.

The Telesophy System

AS AN INTRODUCTION TO THE TECHNOLOGY required to implement an electronic community system, its predecessor will be briefly discussed. Then the specific support for technology and for sociology will be described. Finally, the existing prototype will be discussed, along with future plans.

During the 1960s, Douglas Engelbart's NLS project carried out a pioneering effort to build a tool to "augment the human intellect," a single computer system that enabled a researcher interactively to manipulate all their knowledge [3]. The resulting system could manipulate collections of documents consisting of a hierarchy of paragraphs, which were interconnected on the basis of similar words. There was extensive support for collaboration, such as a shared journal that kept a record of annotations and revisions [4], and support for remote access over ARPANET. The project gathered a

group of devoted users, eventually totaling several hundred, and attempted to support a few small specialized communities [5]. Typical users were information specialists whose jobs involved examining and writing large formal bureaucratic documents with detailed hierarchical structure.

During the 1980s, the present author carried out a project to investigate whether sufficient technology was available to build a complete community system [12, 13]. Much infrastructure had already matured—for example, hardware technology such as large-memory graphics workstations and high-speed fiber networks and software technology such as bibliographic information search and object-oriented programming systems. This project was called *telesophy*, or wisdom at a distance, to indicate that the system was intended to support transparent manipulation of knowledge across computer and communications networks. The concept of telesophy as an all-inclusive network system was intended to be analogous to the concept of telephony, with the ultimate goal of supporting transparent manipulation of “all the world’s knowledge” just as the telephone system supports transparent connection of “all the world’s telephones.”

As part of the Telesophy project, a prototype of a complete community library was constructed, including both data and software [13, 14]. The system forms a distributed digital library, which enables fast browsing across a wide range of data physically distributed across a network. Data are collected from external sources and transformed into uniform objects, which can be manipulated by a uniform set of commands, regardless of the original physical data type. The set of objects is called the *information space* and may be distributed across many machines within a network. Retrieval takes place transparently, regardless of the actual physical location, and is done either by associative search or by following links between objects. The software runs on Sun workstations and has been distributed to more than forty-five sites. The software contains a custom window manager, object system, network handler, and text searcher.

In the main configuration, a wide variety of data was collected and transformed into units in the information space. This collection represented a good sample of what is currently available electronically. It also spanned the range of different media types, from text to graphics to image to video, as a test of the system’s ability to support type transparency. The prototype space consisted of some 300,000 items from some twenty data sources. The informal material were short text messages that included bulletin boards, electronic mail, wire services, and notes. The formal material included bibliographic citations with abstracts covering computing from INSPEC™ and biology from Medline™, and also full text of magazine articles and movie reviews. The pictorial material included line drawing graphics, black and white images, color magazine figures, glossy photographs, and videodisc stills. Finally, the video material included played segments from a variety of educational and entertainment videodiscs. The data thus spanned the range from informal to formal material, as well as including material such as pictures and graphics. A few materials with code were also collected, including playing of videodisc segments and stored queries that are executed on-the-fly to provide a different result each time.

The software supports transparent, distributed information retrieval. Every data item is searchable by combinations of phrases on all the text associated with the item, such

as the abstract, body of text, or picture captions. Searches can be done across all the different databases and all matching items returned. The databases can be physically distributed across machines in a network. In the prototype configuration, the twenty databases were spread across three large file servers connected by a building Ethernet so that any appropriate workstation (i.e., a Sun-3) could access them. The database searching takes place on the file servers, while the user interface and information manipulation take place on the user's workstation. The Telesophy system concentrated on supporting associative search as in information retrieval systems rather than following of links as in NLS or hypertext systems; tradeoffs between search and navigation are discussed in [15].

The software implementation is tuned to support fast browsing. The data are fully indexed so that processing a query typically takes one to two seconds. The resulting items are then downloaded from the remote file server over the network to the local workstation. The interaction is "instantaneous" (less than one second) for displaying and page flipping one-line summaries of query results or zooming into the complete items. This same speed has been maintained for the vast majority of the items (text and line drawings) across a variety of physical networks: building Ethernet, campus Ethernet, and a WAN (wide-area network) consisting of two building Ethernets forty miles apart connected by a private T1 line.

In addition to supporting library browsing, the Telesophy system also supports kinds of community sharing. All of the external data items are represented as information units (IUs), the collection of which forms the information space. There is a single set of commands for basic manipulation of any information unit, independent of type and location. The user can perform an exploratory browsing session, issue multiple queries, and save a collection of selected items as a new information unit that can be indexed and placed back into the space. These collections are a simple form of "metalevel" grouping, of classifying sets of items of different types from different databases into new semantic groupings, which can arise by saving the results of a simple query or from the results of considerable searching and analysis. Since all users, regardless of their physical location, access the same information space, these new composite IUs, which form regions in the space, are automatically shared with other members of the community.

Proving the viability of a new communications medium requires demonstrating that its implementation is technically feasible and its deployment causes a sociological change. The Telesophy system demonstrated the technical feasibility of building a community system, but failed to achieve widespread usage due to the difficulty of obtaining suitable data in electronic form for the needs of the user community—electrical engineers in an industrial research laboratory. Experience with physical libraries has shown that one of their most important features is complete coverage, that is, essentially all materials on the covered subjects are available. Coverage is even more important in an electronic community library since a key feature is rapid annotation of existing material. During the course of the Telesophy project, it became clear that demonstrating the requisite sociological change would require carrying out a large-scale trial with a specialized community, thus prompting the beginnings of the Worm Community System project.

Enabling Technology

THERE ARE A NUMBER OF TECHNOLOGIES REQUIRED to implement an electronic community system effectively. This section discusses one of the most important—the representation for the knowledge in the community library. This builds upon the experience from the Telesophy system.

Information Spaces

The data model for a community system must support uniform commands for browsing and sharing across the complete spectrum of community knowledge. This requires supporting features not well supported by the models underlying traditional database management systems. Community knowledge spans a wide range of types, each requiring its own operations for search and display. Community knowledge is interconnected and needs an efficient representation for making relationship links between items. Community knowledge exists across many sources, which can be distributed across a network. The relational model, for example, cannot easily support multiple types or arbitrary links between arbitrary groupings. An appropriate object-oriented model can, since each type of object can have its own set of operations and each object can have its own set of pointers to other objects. A community system uses a particular kind of federated heterogeneous distributed object-oriented database, called an *information space*.

Information spaces support uniform manipulation of heterogeneous data items by transforming them into homogeneous information units. The generation of an information space begins with data already existing in some external source. The format of this data is administratively transformed into a canonical internal representation called an *information unit*, or *IU*. An information unit is an encapsulated object, in the sense of an object-oriented programming language, which has an associated set of operations to provide manipulation capability for its particular data type. Every “database” thus has a set of transformation routines and every “data type” has a set of data operations. Once the data items have become information units, there are a set of generic operations available for performing on them. These generic operations support uniform commands at the user level for such functions as search, display, and grouping. Thus, a user of an information space need only learn one set of commands to manipulate information units, which operate uniformly across a wide range of external data types. Each information unit may be connected to other units to represent a semantic relationship and collections of information units may be grouped into new composite units. An *information space* is a set of information units and their connections. Logically, it is a single uniform graph structure, although physically it may be composed of many different sources of data of many different types stored on many different machines in many different locations spread across a network.

There are several levels of representation in an information space. Data exist in the external sources and are transformed into information within the space. Knowledge, in the sense of *community knowledge*, is represented by the different components of

information units. Any IU can be annotated; a typical *annotation* is a note stating some additional feature of the encapsulated data, for example, this gene may encode this function. Any two IUs can have a relationship specified between them; a typical *connection* is a link to another IU supplying additional information, e.g. this article discusses this gene. Any collection of IUs can be grouped into a single composite IU which forms a region in the information space; a typical *region* is a set of IUs on the same topic, for example, all genes coding for mechanosensory deficiencies. Since every IU has a unique identification within the entire space, it is possible to implement a uniform mechanism for forging and maintaining these groupings, even across sources. As discussed below, every IU also has specification to provide publication control over the sharing of these groupings.

Forming the Space

Anything accessible may potentially be incorporated into the space. That is, all data reachable via the underlying network for which appropriate transformation routines exist can reside logically within the information space. When administration is done to bring data physically into the space depends on ease of reliability and maintenance. In many cases, maintaining the data directly in an external database is the most convenient; in this case, data items are transformed into information units only when they are actually retrieved (and then only temporarily during use so that any updates must be written back into the database itself). If the data are to be maintained directly in the information space itself, the data items are transformed once into information units when they are brought into the space and then any updates are performed by operations within the space. Since maintaining consistency and correctness of large amounts of data requires considerable system support, initial implementations of information spaces will likely rely on existing database management systems to provide maintenance, transforming external data items into internal information units on-the-fly or periodically whenever the database has been significantly updated.

In the worm information space, for example, there are a variety of methods for incorporating external data and software. The support for these may be handled internally (as objects brought into the system) or externally (as objects existing outside the system). Some external data are read in from text files, then handled by internal software. For example, the gene list is a text description kept in a file, then supported by the built-in text display. Some external software is invoked as a separate process with arguments. For example, the sequence map display is called as an external program. Some external software is invoked with objects passed in and out. For example, a sequence analysis program is passed sequences in a canonical textual format and returns text that is transformed back into sequence objects. Finally, some external software supports its own classes which are directly communicated with, providing internal software with direct interactive access to external objects. For example, the genetic map displayer is an external program that implements an annotation command that invokes the internal support for annotating the objects belonging to the external program.

The major generic operations built into the system, as part of the IU class definition, are the support for grouping. These include connection links and region sets. Other operations, which provide support for the uniform user commands, are implemented at the individual subclass level, for example, those for search and display. This enables the system to support many different types of search (e.g., text and sequences), and of display (e.g., text and maps). Some of the type classes are available in essentially every community—for example, an atomic class for text and a composite class for some kind of hyperdocument. Other types are specific to individual communities—for example, an atomic class for gene and a composite class for genetic map containing gene positions. The object structure of information units enables an electronic community system to be extensible, with a base set of classes that can be augmented by specific classes for a specific community.

Enabling Sociology

THE ABOVE DISCUSSIONS HAVE INDICATED THAT IT IS TECHNICALLY POSSIBLE TO collect a significant amount of community knowledge and make this easily available to community members. Ensuring their active participation in this electronic experiment requires resolution of the following sociological problems, among others.

Editorial and Quality Control

Published literature typically goes through a careful refereeing process. This is also true of the archival data, where there is typically a trusted central administrator who performs editorial quality control. With informal information or unpublished data, especially when entered by the users, quality control becomes significant.

The solution to quality control in the printed literature is to have a range of editorial review that leads to a spectrum of documents ranging from lab notes to working documents to internal memoranda to newsletter announcements to conference papers to journal articles to research monographs to textbooks. A similar spectrum has emerged in electronic bulletin boards. In public boards, anyone can post any message. In moderated boards, all messages go first to a moderator who eliminates those that are on wrong topics, redundant, or inflammatory. In edited boards, the editor passes judgment not only on topic but also on quality and format. There has been talk of true refereed boards with long articles, but not many examples exist.

A community system should provide a mechanism for "levels of editorial release," that is, how carefully checked an item is before it is released to the community. Following on the experience of electronic bulletin boards, the spectrum of editorial control should include: posted, moderated, edited, refereed. The system does not, however, determine the policy of which level an author chooses for a particular item or who performs the function of the editor for which items. An appropriate set of conventions will have to evolve for the electronic community library, just as such a set of conventions has already evolved for electronic bulletin boards. Based on experience with the worm community in the past, editors will emerge who can provide

appropriate levels of quality control for each data source and who are sufficiently respected by the community so that their blessing of the data is trusted.

The level of editorship should be recorded on each item in the information space, because this is of interest to the researchers who are evaluating the suitability of particular information units for their current purposes. This is a form of policy that permits the individual users to choose for themselves whether they are currently interested in refereed facts and data or in rumors and notes.

Privacy and Reward Considerations

Another problem in extending the community library beyond formal data is whether the members are willing to share the data before it has been formally published. The tradition of freely sharing unpublished data is a primary reason for choosing the worm community for the initial experiment. But there is a significant problem in any scientific community for establishing credit and priority, particularly as competition becomes more intense.

The community system should provide the mechanism of "levels of privacy release," that is, who is permitted to view which material. Sample levels include: private (user only), colleagues (local), colleagues (global), community. As with the editorial release, the policy for each item is individually determined by the author and can be changed as the item evolves in maturity and quality. Each researcher can also determine who is permitted to view each level of release, that is, who their colleagues are.

Conversely, for searching purposes, each researcher can use the privacy level to help determine the appropriateness of those items in the information space that they have permission to access. It should be noted that the privacy levels enable the community system to support services equivalent to electronic mail and bulletin boards.

An issue related to privacy is rewards. What reward does the author of an information unit receive? It will be a long time before the prestige of making a connection in information space rivals that of publishing a paper in a journal. The system can provide the mechanism of a super citation index, by keeping track of the frequency that an item is retrieved and the number of times a connection is made to it. Hopefully, these usage statistics will aid in establishing policies for electronic publishing. Also note that every information unit has complete attribution of its creation, that is, author and date. In fast-moving fields with extensive electronic coverage, this could provide a method for establishing priority and credit.

Prototype Community Systems

THE FIRST RELEASE OF THE WORM COMMUNITY SYSTEM was completed during the summer of 1991 and is now in the labs of the initial test users, who are using it to browse the data and beginning to add annotations.

The current community knowledge spans the potential range. It includes fairly complete archival data, such as the list of gene descriptions, the genetic map, the physical map, and many DNA sequences. It includes abstracts of most of the archival

worm literature. The worm newsletter has been completely scanned and the text recognized, so that articles can be searched for, then displayed as formatted text with accompanying images for the figures. Unpublished data are available, such as standard strains and a person directory, plus a sampling of other data from individuals.

The software functionality also spans the potential range. Searches can be done across all the sources for text phrases. An extensive set of links has been made between information units by a variety of automatic and manual means. These links can be followed from any IU to the related set of IUs. Sets of IUs can be selected and grouped. Several external analysis programs can be called to provide displays of worm data for the genetic map and the sequence coding map. Finally, an annotation facility is available which permits a note to be added to a set of IUs giving additional information about the group. This note may include embedded links as well as text. When saved, annotations are released into the information space, where they can be manipulated as ordinary information units.

Sample Session

Figure 2 is a screendump from a sample session with the Worm Community System. This session is a summary of the interaction with a biologist, interested in sensory neurons, who is attempting to discover which genes in *C. elegans* control the sense of touch, mechanosensation. The information space enables the biologist rapidly to locate all such known genes and retrieve information about them.

The session starts with the user entering a search for the keyword "sensory" as shown in the topmost Search Control window. The search is performed across all objects of any type contained in the information space; the number of objects is shown in the upper right. The window below Search Control labeled Search: "sensory" contains a summary of the set of objects in the worm information space matching this keyword (containing that text string). Each object has a one-line summary (uniform for all types) that can be zoomed into by pointing with the mouse and double-clicking. The selected object is displayed in the bottom window. It is a literature object containing a citation and abstract from the journal articles about the worm.

In addition to associative search, units in the information space are interconnected and the user can follow links to navigate to related units. For the worm space, literature objects are linked to all genes described in the article. Figure 3 shows a link following. In the window labeled Search: "Traversal Set," the user has requested all objects linked to the selected literature object and the system displays one-line summaries of the set of matching gene objects. The user selects one gene "mec-3" and zooms into its description, which is displayed in the bottom window. This description shows that mutations in the gene indeed make the worm insensitive to touch.

The user now wants to see where the gene is located physically on the DNA of the worm and issues the "show physical map" command on the selected gene in the Traversal Set. Figure 4 displays a section of the physical map of the chromosome showing the known locations for a variety of cloned DNA fragments. The window labeled Contig #423 displays the region containing mec-3. This window is not just a




Figure 2. Search





Figure 3. Navigation



line drawing but a live first-class graphical display of a composite object "physical map" that contains many subobjects of type "clone." Thus, the individual objects can be manipulated and interacted with. In this session, the user zooms into the clone on the map containing *mec-3* and displays its DNA sequence in the window at the bottom of the screen. An external analysis program might now be invoked (but not yet in this prototype) to compare this gene controlling touch in worms to a library of the sequences for genes in all organisms to identify similar genes in humans.

Finally, the user checks whether other community members have added any informal information about the gene of interest. Figure 5 shows the result of zooming into the physical map entry for *mec-3*. Three IUs are displayed, corresponding to the gene, the clone, and the DNA sequence. The gene has a checkmark by its summary line indicating that an annotation is available. Issuing “show annotations” brings up the window discussing “Touch Receptors.” This has a number of embedded references denoted by %-. Zooming into the reference stating that *mec-3* is a homeobox retrieves the paper in the bottom window. The user has thus made good use of the value-added informal knowledge to find a relevant specific paper concerning the gene of interest.

Future Plans

The current prototype system is written in GNU C++ and runs under the UnixTM operating system, typically on a Sun SPARCstationTM. The external sources are maintained in files in text form and transformed into information unit objects when




Figure 5. Annotation

the system starts up. All the software is custom written, including the object manipulation. The current system runs all in memory and takes about 11 megabytes when loaded. This comprises some 18,000 objects, the bulk of which are physical map entries, while the bulk of the size is literature items. Some of the test sites run the system on a local Sun workstation directly. Since the display uses X-windows, others run it remotely on an Apple Macintosh™ II running MacX™, with acceptable response over a local area network.

This first release is now in the labs of the initial users, on the order of ten laboratories. Initially the goal is to recruit enough users to support a "fair test" of this kind of system. These users must be enthusiastic enough to use a preliminary system and influential enough to have their reactions taken seriously by community members. In addition, geographical distribution is important since the experiment is a test of a nationwide electronic scientific community.

The feedback from this release is being used to design the second release. This version will be a distributed system with separate modules for database searching, for object manipulation, and for user interface. Computer scientists associated with the project will be experimenting with what caching and protocol technology is necessary to provide interactive retrieval across the NSFNET. The plan is for this version to be a fully featured system that is propagated to a significant fraction of the worm community. Sociologists associated with the project will be investigating its usage to understand the effects on the community's communication patterns.

In the longer term, as the system for supporting the worm community becomes functional, the software will be made available to other communities. The next

electronic communities will probably be molecular biologists whose communities are also organized around experimental organisms, including the bacterium *E. coli*, the fruit fly *Drosophila*, the weed *Arabidopsis*, the slime-mold *Dictyostelium*, the alga *Chlamydomonas*, yeast, mouse, and man. The problem for many of these communities will be the lack of coverage of available data in electronic form, but much of the software should prove transferable.

As attempts to build more electronic community systems begin, more will become known about the characterization of communities. Many factors play a role in the suitability and usefulness of such a system to a community. These include: data (extent of coverage and vitalness of need); maturity (size and age of the community); competitiveness (readiness to share, pace and stakes); sophistication (computer literacy, tolerance for new technology); and many others. It may eventually prove possible to tailor an electronic community system to be more effective for a given set of community characteristics.

Toward Electronic Systems for Organizational Memory

AN ORGANIZATION IS IN MANY RESPECTS SIMILAR to a community. It consists of people with common interests and shared values. So the knowledge in an organization is similar to the knowledge in a community. This knowledge might be termed *organizational memory*, that is, the knowledge that enables the organization to continue to function effectively. This is the permanent knowledge, as opposed to the transient knowledge generated during meetings. As with communities, organizational memory includes not just the tangible information in designs and memoranda, but also the intangible information in company procedures and values. The permanent memory in an organization that can enable it to outlive its founders is contained in both the company products and the company culture. Recording this memory and making it easily accessible electronically would clearly be of enormous use to the functioning of organizations.

The knowledge-encoding and software-system techniques developed for an electronic community system will likely be relevant to building electronic systems for organizational memory. An industrial organization, for example, has similar knowledge to the scientific community described in this paper. There are archival data, such as design specifications and product evaluations, and intermediate data, such as test results and market surveys. There is formal literature, such as technical memoranda and progress reports, and informal literature, such as design notes and meeting minutes. An organization also has similar needs to manipulate this knowledge. There is a need to browse the knowledge, to filter out selections relevant to the problem, then to share annotations of these selections with other members of the organization.

The sociology of an organization tends to be somewhat more rigid than a scientific community. Thus, although the faster pace will likely require the fast distribution of the knowledge, greater control over its dissemination is likely to be important. A finer granularity of specification and tracking of the editorial and privacy controls is likely to be necessary. For example, a design must be approved by a specified series of people

and product information is available only on a need-to-know basis. There may be other controls to support the policies and procedures, and to regulate the flow of information within the organizational structure. Finally, there may be other types of functionality necessary to capture some degree of the company culture. For example, there may be precise style and content constraints on hyperdocuments in the organization's information space.

Large-scale experiments in real organizations will be necessary to assess the value of electronic community systems for supporting organizational memory. Preliminary evidence indicates that this technology will be valuable to the scientific community, so great potential exists for its value to the business community as well.

REFERENCES

1. Coulson, A.; Sulston, J.; Brenner, S.; and Karn, J. Towards a physical map of the genome of the nematode *Caenorhabditis elegans*. *Proceedings of the National Academy of Science*, 83 (1986), 7821-7825.
2. D. Crane. *Invisible Colleges: Diffusion of Knowledge in Scientific Communities*. Chicago: University of Chicago Press, 1972.
3. Engelbart, D., and English, W. A research center for augmenting the human intellect. *Proceedings of the AFIPS Fall Joint Computer Conference*, 33 (1968), 395-410.
4. Engelbart, D. Collaboration support provisions in AUGMENT. *Proceedings of the AFIPS Office Automation Conference*, Los Angeles, February 1984, pp. 51-58.
5. Engelbart, D. Coordinated information services for a discipline- or mission-oriented community. In *Computer Communications Networks*, R. Grimsdale, ed. NATO Series, vol. 4, Noordhoff, 1975, pp. 89-99.
6. Licklider, J. *Libraries of the Future*. Cambridge, MA: MIT Press, 1965.
7. Licklider, J.; Taylor, R.; and Herbert, E. The computer as a communication device. *Science and Technology* (April 1968), 21-31.
8. National Science Foundation. Towards a national collaboratory. Report of an Invitational Workshop, March 1989, Directorate for Computer and Information Science and Engineering.
9. Quarterman, J. *The Matrix: Computer Networks and Conferencing Systems Worldwide*. Englewood Cliffs, NJ: Digital Press/Prentice-Hall, 1989.
10. Reid, B. The USENET cookbook—an experiment in electronic publishing. *Electronic Publishing*, 1 (1988), 55-76.
11. Roberts, L. The Worm Project. *Science*, 248 (June 15, 1990), 1310-1313.
12. Schatz, B. *Telesophy*. Technical Memorandum TM-ARH-002487, Bell Communications Research, August 1985.
13. Schatz, B. Telesophy: a system for manipulating the knowledge of a community. *Proceedings of IEEE Globecom '87*, Tokyo, November 1987, pp. 1181-1186.
14. Schatz, B. A prototype information environment. *Proceedings of the 2nd IEEE Workshop on Workstation Operating Systems*, Pacific Grove, CA, September 1989, pp. 118-124.
15. Schatz, B. Searching in a hyperlibrary. *Proceedings of the 5th IEEE International Conference on Data Engineering*, Los Angeles, February 1989, pp. 188-197.
16. UNESCO. *UNISIST: Study Report of the Feasibility of a World Science Information System*. Paris: United Nations Educational, Scientific, and Cultural Organization, 1971.
17. Wood, W., ed. *The Nematode Caenorhabditis elegans*. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press, 1988.

Copyright of Journal of Management Information Systems is the property of M.E. Sharpe Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.