

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

**Management of the Design Process:
The Impact of Information Modeling**

**J. Robertson, E. Subrahmanian,
M. Thomas, A.W. Westerberg**

EDRC 06-179-94

MANAGEMENT OF THE DESIGN PROCESS: THE IMPACT OF INFORMATION MODELING

Jerry L. Robertson

Consultant, 20 Westwind, Sand Springs, OK 74063

Eswaran Subrahmanian

Engineering Design Research Center

Carnegie Mellon University, Pittsburgh, PA 15213

Mark E Thomas and Arthur W. Westerberg

Dept of Chemical Engineering and Engineering Design Research Center

Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

An information revolution is impacting how companies will operate in the future. Companies are looking for models that they can use to improve organizational efficiency. Early models decomposed organizations into workers and management and ascribed roles to each; newer models such as that by Denting concern themselves with integration of the parts into the context of the outside world. We suggest an analogy to decomposing processes into unit operations which must then be integrated into total processes within the world market, suggesting that, while the words are different, the process design activities may be very similar. Other important chemical engineering concepts such as process control and information transfer appear to be applicable to organizational management as well.

After examining management processes and how information flows in a company, we explore different technologies for information management/consolidation, noting how well each permits the gathering, structuring and using of the data both to operate a company and to understand and improve its performance, particularly in its ability to carry out a complex process such as design. Particular attention is paid to databases, the World Wide Web, and a new information management system under development at Carnegie Mellon University called n-dim.

The paper ends with a brief view into Etopia, an information society some time far into the future.

Introduction

This paper speculates about how changes in work processes and the way information flows in companies along with new developments in information management techniques will create opportunities for improving the design process. It describes how models of the design and operation of organizations seem to parallel models of chemical engineering processes and suggests that similar information management techniques will apply to both. It describes some of the developments in information technology that will apply to both.

Trends in Corporate Organization

Major transformations are underway in management literature and practice. Almost all institutions are reshaping their patterns for organizing work and the work place. They are scrutinizing traditional folklore about management truths and testing new management

ideas, while also examining concepts about managing

for and achieving? <!"** P ^ f : We *!'.*.*.* **
the activities of institutions, in their current search for improved work processes and organizational design, parallel our activities aimed at improved chemical and petroleum process design and that computer-based information technology will have major impact on both.

Breaking Down Work Efforts—Decomposition

Taylor's [1911] classic of management science literature has strongly influenced the configuration and activities of the workplace. Taylor proposed to improve the efficiency of the organization by substitution of science for individual judgment. There was only "one best method"⁴ which guaranteed maximum efficiency. The worker was considered to be part of the machine. It was, therefore, unnecessary for a worker to understand why or exactly how the system performed. Explicit job descriptions, time and motion studies, and work method standards were expected to assure work efficiency and free the worker from the drudgery of work practices. Management's functions in his model are to enforce work standards and provide the best tools and working conditions. It provides the understanding of the overall goal and the means to achieve it. As organizations grew in size and complexity in the early 20th century, this model enabled great increases in efficiency. It remains central to many organizations today.

Drucker's [1973] view of management is broader. "Management has to give direction to the institution that it manages."⁹ This includes establishing mission, setting objectives, and providing resources for the institution. Management skills include: "communications within organizations; making decisions under conditions of uncertainty; and strategic planning." The importance of the customer is emphasized: "There is only one valid definition of business purpose: *to create a customer.*"

We suggest that Taylor's view of organizational effectiveness and Drucker's view of management tasks might parallel chemical engineers' view of their domain in terms of unit operations and unit processes. Decomposing the process systems into their units and understanding the detailed functioning of these units enables the understanding of the overall process just as decomposing organizational work functions enabled improving organizational efficiency.

Analysis of Work Activities

In the early 1980s many observers viewed US industry as being in organizational crisis. Peters and Waterman [1982] searched for excellence in America's best run corporations and found, among other things, three major characteristics: 1) that customers reign supreme, 2) that a high level of employee dedication and enthusiasm was essential and 3) that trial and error was not only accepted but encouraged. The late W. Edwards Deming taught his 14 Essential Points for Managers [Walton, 1986; Scherckebach, 1986]. The Deming method emphasis is also on the customer as well as improving consistency, instituting leadership, eliminating employee fear, and training and retraining. Some essential elements of Total Quality Management (TQM) include: 1) an accurate and complete understanding and view of the

customer; 2) a need to understand the entire work process; 3) systematic measurement of performance to assure minimum variation; and 4) the need to continuously improve the work performing system.

Several paths are being used to achieve TQM including assimilating the Deming steps, utilizing the Baldrige Award criteria as measurement techniques, and undertaking ISO 9000 certification. In many ways these paths are similar. An advantage of the Baldrige Award criteria is that it provides metrics for measurement of the state of the organization and, therefore, also can be used to indicate improvement quantitatively. While it can be argued that these metrics are arbitrary, they nonetheless provide valuable indicators to management.

The ISO 9000 series is a set of five quality system standards which include documentation of activities and procedures, guidelines for management reviews both in frequency and content, provisions for organizational self review, and training. Rather than identifying and distinguishing measurable characteristics as found in most technical standards, the ISO 9000 series provides management guidelines and models that are aimed at assuring customers that their expectations will be met. They are characterized by "considerations" rather than "directives".

ISO 9000 certification is becoming essential for business, particularly in the European Community [Corrigan, 1994]. Customer requirements were cited as the leading reason for registration by 26% of 110 chemical companies in a recent survey. ISO 9000 registration reached about 3,600 in the first quarter of 1994 compared to 1,500 a year ago and 400 two years ago [Thayer, 1994].

We suggest that if the Taylor/Drucker methods for analysis of management is parallel to chemical engineering unit operations, the Deming/TQM approach is analogous to overall process descriptions—the integration of parts and their context with the outside world. Analysis of the overall process as an entity and examination of the interactions of the components to determine how to improve the process appear to be characteristics of both. Using techniques that prove successful in management may be beneficial for improving process design methods.

Synthesis of New Work Processes

More recently, the impact of the information age and the necessity of examining the details and objectives of the entire work process is being described. Naisbitt [1985] described the transformation of organizations and specific work activities that were beginning to emerge as a result of the information age. Flatter organizations with less bureaucracy, networks where everyone learns from everyone, where intuition becomes valuable because there is so much data, where much of the work is performed by contract rather than hired staff are part of the "Re-invented Corporation". Middle management is

replaced by the computer and people managers by independent, competent and self confident, self managers.

'The **fundamental** rethinking and **radical** redesign of **business processes** to achieve **dramatic** improvements in critical contemporary measures of performance...*' is how Hammer [1993] describes the activity of *Re-engineering the Corporation*. He emphasizes the difference between incremental continuous improvement and fundamental change to the work processes, organization and culture. The business process is the complete action of creating and delivering the desired products to the customer and must be understood in its entirety. In contrast to Taylor's work breakdown, Hammer's emphasis is on work integration. Activities of organizational entities must bring added value to the core business as a whole rather than be justified on an internal unit goal. Even inter-company organizational boundaries are crossed. Information technology is a key enabler for the re-engineered corporation. It does this by creating opportunities to change the work process dramatically and not just doing the same thing faster. Radical redesign of the complete organization from top to bottom and a dramatic change in the way of doing business is also cited by the National Research Council [1991] as an ingredient for Designing for Competitive Advantage.

Work and Chemical Processes-the Future

We suggest that the most recent approach to organizational management analysis strongly parallels the state of chemical engineering process design analysis when process synthesis was embryonic. We see parallels between Hammer's work and retrofit design. Institutional management may benefit from the systematic techniques developed and proven for process synthesis. Furthermore, the needs of the process design community and the general business community in terms of information consolidation and management appear to be strongly congruent. The challenge will be to integrate developments in the common elements of these communities which have diverse histories, cultures and even language. It is apparent that one of the certainties of the future is uncertainty and that changes will occur not only in information technology, but in all aspects of business. The Learning Organization as advocated by Senge (1990) will be better able to respond to these changes. "Total Systems Thinking," seeing the whole ("The Fifth Discipline"), is a major attribute of this organization. It should have the same characteristics as

a self-adaptive organism, adjusting to changes in the business environment and technology. System archetypes [Senge, 1990] which help understanding of business detail and dynamic complexity and which can be connected to form different business structures might be a useful starting point for business process synthesis. Furthermore, concepts in chemical engineering process control might be useful in the analysis of business environment and organizational dynamics.

The preceding review focused on elements common to chemical engineering and business processes. A basic difference is the human element in business processes. Morale, participation with high performance teams, innovation, motivation to learn, and just plain thinking hard are human characteristics for which information technology has little to offer at present. Both business and process design success relies heavily on these characteristics.




Fig. 1 Total corporate information flow loop for a typical chemical or petroleum industry institution

The Corporate Information Loop

As we discuss later, information consolidation technology is evolving rapidly. While the design activity represents only a small fraction of total resources involved or total information flow in any corporation, it is strongly influenced by changes in information flow. Fig. 1 shows one perception of a total corporate information flow loop for a typical chemical or petroleum industry institution. All commercial businesses utilize work processes to convert some input to some output. Each of the nodes

represents a subdivision of the overall organization. In this simple example/the Procuring node represents obtaining raw material or other inputs; Manufacturing represents the transformation process and Marketing represents dialogue with the customer. Corporate Management makes up the final of the four major nodes. Further sub-nodes depict major activities within each of the major divisions. It is likely that there are further sub divisions in most organizations. (For example, offsites design as a sub division of design and cooling water system design as a sub division of offsites design).

There are three points to make with this diagram: 1) almost everything done within the design sub-node is influenced by some information from one or more of the other sub-nodes, 2) organizational control, as usually practiced, requires that information based on current data be approved by the management of each of the sub-nodes before it is shared with other nodes and 3) all information is not shared among all nodes. As information becomes more broadly shared and organizational culture changes, it is likely that the information available for design activities will also change. It appears likely that decisions that were made by parts of the organization and documented by providing "functional or job specifications" are more likely to be made by the design team. For example, data about costs of equipment, catalysts, utilities, etc. will be available in "real time". Product quality requirements and projected sales data would also be current as would knowledge about the availability financial resources. This activity would have such broad corporate ramifications that the team might be referred to as the "future team" rather than the design team.

Information

We are in the age of Information Technology. But what is information technology? Let us first define information. One often hears discussions about the differences among data, information and knowledge. Typically this discussion proposes that *data* is unprocessed, plentiful and likely to contain both meaningful and irrelevant things. An example is all the readings from a set of sensors for a plant. When one takes lots of data, removes the irrelevant parts, discovers the hidden messages in it and organizes it to expose patterns, one is transforming data into *information*. When one converts information into insights that aid in one's work, one is transforming it into *knowledge*. These steps are abstraction processes in that lots of data are summarized to form a more useful and much more compact way to describe it. An example of abstraction is to take a parts list and abstract it by saying this list describes a distillation column, an interpretation that immediately gives considerable understanding of it to a chemical engineer. It becomes evident that one person's knowledge can be someone else's data. These definitions

are, therefore, relative to one's perspective. In what follows we shall be concerned with aiding this abstraction process.

Information Management/Consolidation Technology

Information management/consolidation involves the discovering and/or creating, the abstracting, the interpreting, the organizing, the using, and the sharing of information. The computer, and very specifically the networking of computers, is at the heart and will continue to be at the heart of all new information management/consolidation technology.

Consolidation technology involves the use of analysis tools to abstract information content out of data. An example would be to input the data from a set of process sensors into a computer model which then determines the material and heat flows for that process. Another would be to take noisy data and smooth it using digital filtering techniques.

While such analysis tools are extremely important and worthy of a lengthy discussion, we shall concentrate in this section what we call information management (IM). We shall, in particular, emphasize ways to discover information located elsewhere, to organize information for the purpose of exposing relationships among it and for sharing it both directly and by making it available for others to discover.

A simple form of managing information is to organize files containing it under a directory/subdirectory tree as is done in all computer operating systems. One typically organizes such a directory system to put things which are related into the same subdirectory. Through the use of "links" (a special pointer to a file), a file can appear to be in more than a single subdirectory. If one browses a subdirectory, the operating system lists a link exactly as it lists a file name. Clicking on a link will move one to the subdirectory containing the file and will open the file. The operating system will now be pointing in the subdirectory where the file is located with the disadvantage that it (and often the user) will not remember how it got there. Files can contain any of a variety of things: text, drawings, audio recordings, film strips, and so forth.

A second way to organize information is to place it in a database, and a related third way is a document management system. We shall assume the reader is familiar with the concepts related to both of these technologies.

A fourth approach is to organize information into a hypertext. A *hypertext* is a text or graphics document with buttons attached to it. Clicking a mouse on one of these buttons causes a hypertext system to open up another page of text that the author has *linked* with the current document by associating it with the button. On-line help systems today often are hypertexts. The author places buttons, which are often invisible, under a piece

of text (or within a drawing) so one gets the impression that one is clicking on the text. Doing so will bring up a page relevant to that piece of text. A hypertext is, therefore, a complex interlinked document which one can read by following several different paths through it. The hypertext concept is an important ingredient of the world wide web (WWW).

Placing meaningful labels on the links which connect two documents in a form the IM system recognizes extends the capability to relate information ever further and is part of the semantic network modeling one often uses to create data models for databases. Such labeling is also an important ingredient of our information modeling environment n-dim (n-dimensional information modeling).

There are many properties we hypothesize we want to have in IM systems. We want to keep and link information on people, on their activities, on the formal and informal organizational structures relating those people, on the numerical and symbolic data describing physical artifacts such as the data for a flowsheet computation, on the group development of arguments such as described by gIBIS [Conklin and Begman, 1988], on drawings, on correspondence and e-mail files, and so forth. We want to be able to store tremendous amount of information. We want to be able to find things in this information by browsing, by formal searching and by being pointed to it because someone else or the IM system concludes it might be interesting to us. We usually want very fast responses. We will often want the IM system to record enough history of the data creation process that we can study the process and/or reproduce the arguments behind the *why* of the data and not just values of data. The argument that the annotated blueprint for a design is not enough supports such capture. Often, especially for a new activity, deciding what information to gather, how to structure it and how to share it is a major part of the effort one makes in managing it. We want these systems to support exploration among alternatives for organizing and interrelating it.

IM cannot be solved through technology alone. With more information kept, a company may expose itself to litigation where the evidence needed by the other party will likely be sitting there in the IM system. There are also significant "people" issues one must understand also when developing and using this type of technology. For example the maintenance of a more complete history may make a participant very nervous about how his/her boss will use it when the annual review comes around.

We shall now examine the World Wide Web (WWW) and /i-dim in more detail. We shall argue then how databases, the WWW and /i-dim are each useful for information modeling, assessing both advantages and shortcomings.

The World Wide Web

The World Wide Web is a system that is supporting the creation and sharing of a global hypertext web of information. It physically exists now and is growing very rapidly, both in the information it contains and in the facilities available over it.

Two tools provide access to the WWW. The first displays any document written in the HyperText Markup Language, HTML. An HTML file is an ASCII text file with instructions mixed among the text telling how to display it. The approach is very similar to having a text document in Latex [Lampoit4986] or Scribe [Unilogic, 1985]. The second tool is the locator tool and is one that, when given the address of an object in the form of a Uniform Resource Locator (URL), will automatically retrieve it from across the Internet using its own HyperText Transport Protocols, HTTP. (These protocols are similar to the commonly used File Transfer Protocols (FTP) and to those underlying X-windows.)

A very common program that people use to interface these WWW access tools is Mosaic. One can also interface than using the Emacs editor, for example. Anyone with Internet connections can access the World Wide Web by installing Mosaic (see the box below on how to retrieve the Mosaic interface software).

Getting the Mosaic WWW interface program: Anyone on the Internet can obtain an executable and/or source code needed to run Mosaic. Getting the code is via anonymous FTP. At present support is available for several workstations but not all. FTP to the Internet address *ftp.ncsa.uiuc.edu*. Use *anonymous* as the response to *name* and follow instructions. The binaries are in the subdirectory Mosaic/Mosaic-binaries. They have been compressed using the *gzip* program which is available by anonymous FTP to *prep.ai.mit.edu*. *gzip* is in the subdirectory */pub/gnu*.

When one starts the Mosaic interface, the locator tool retrieves and displays, using the display tool, a designated HTML file called a "home page." This display will contain text intermixed with graphics. Within the text are highlighted phrases (underlined and blue). Each indicates the existence of a link to another computer file elsewhere in the World Wide Web. As one places the mouse arrow (without clicking) over a highlighted phrase, the Mosaic interface displays the associated file address (its URL) in a small window at the bottom of the main Mosaic window. Clicking on the phrase invokes the locator tool which then retrieves the object and opens it, or, if it is a program, launches it (i.e., causes it to start executing).

This object may be another page, a graphics object such as the latest US infrared weather map or an executable such as the program on the library computer

at Carnegie Mellon University that allows one to search CMU's database of article titles, authors, keywords and abstracts. The locator tool is capable of deciding the kind of file it has retrieved and launching it if it is an executable or starting a local software package that can display it, if that software exists; if not available, one gets a simple apology and no display.

A person can write his/her own pages in HTML using a text editor and add him/herself to the WWW at any time.

n-dim

n-dim is an information modeling environment on which the three authors of the this article from Carnegie Mellon University, along with others in the *n-dim* group, have been working for the last four years. From observing different design projects (new control system for power generation with Westinghouse, a connector design with three other Engineering Research Centers and a Database design with Alcoa), we noticed the importance of managing the information that the different groups created and the difficulty all the project participants had in sharing it and establishing a shared understanding of it. Two subsequent projects, one with ABB and one with Union Switch and Signal, have provided us with more empirical evidence of the importance of handling information within a design activity. There are also several other projects reported in the literature that come to the same conclusions [Schmidt, 1993, Grønbaek, et. al., 1993].

To aid the handling of information, we hypothesized that a design team was always *modeling* things. Team members are continually sketching and revising models of the organization, of the process they will undertake - perhaps as an activity diagram, as an equational model of the physical artifact they wish to design, etc. We proposed a general purpose *modeling* environment where one can interrelate all types of models. We also wanted to support activities where the design team needs to learn what information to gather, how to structure it and how to share it. Finally, like the Internet, we argued that *n-dim* has to support users at their desks while they are geographically apart

The environment has to record the process occurring more precisely than anyone can do at present, annotating the mistakes or the reasoning used as well as the final results. It should allow a future team to use this recording as a starting point for a similar design. It should also allow the company to have the information needed to improve the work process itself (i.e., the work process is also a designable artifact).

Recording also suggests big brother sitting there watching (including governmental agencies who demand information). Recording is needed. However, one cannot ignore the social concerns discussed here; we must consider them along with the technical solutions.

n-dim has progressed through three generations of development from a Macintosh hypertext version to a demonstration multi-machine, multi-user IM system. The current version is a complete rewrite, comprising the BOS (Basic Object System) in C and supported by an interactive language called STITCH. We can implement new ideas very rapidly and model hundreds of thousands of objects.

Information modeling: *n-dim* uses nodes and links models. Fig. 2 illustrates. Consider first only the files and the links; ignore the grouping. Labeled links as we show here are not possible in organizing a typical file system. The labels can tell us how we think the files are interrelated. This type of structuring must be supplemented with grouping so we can appreciate which "Data file" leads, for example, to which "Printout". The same file - for example, "My code for X" - should exist in several groups. Fig. 3 shows actual groups typical of models in *n-dim*. They appear as boxes within boxes for the has-part link (the upper box has two files which are part of it: "Object 1" and "My code for X") and with explicit display of any labeled links, here the "annotates" link.

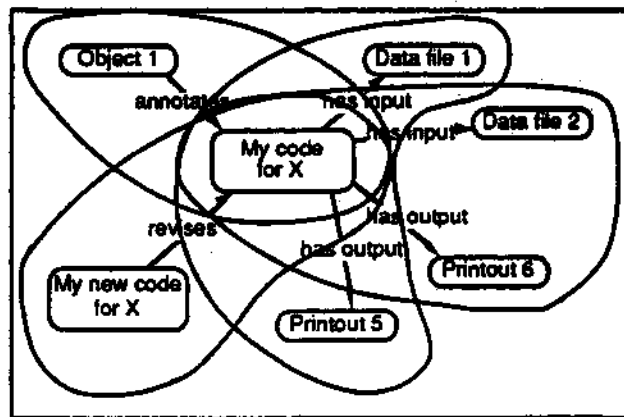


Fig. 2 Grouping files into different views

Atomic model object types in *ft-dim* include integers, reals, complex numbers, frames, file pointers and the like. From these and other previously generated models we construct new *n-dim* models, each of which is a list of links - pure and simple. Links are objects that point from one object to another. In *ft-dim* they always have an associated label, *ft-dim* contains special "has-part" links. Each points from the identifier of the model itself to the identifier of an object that model "contains." It is the display of these objects that gives one the impression that the parts are "in" the model. For convenience, we will often speak of an object being in an *ft-dim* model when in fact the *ft-dim* model only contains a has-part link from the model identifier to the object identifier. Semantically, has-part links are the links one creates when organizing a set of files into a

directory/subdirectory tree or in creating a hypertext. The parts of a directory are its contained files and subdirectories. The parts for a hypertext are those objects associated with its buttons.

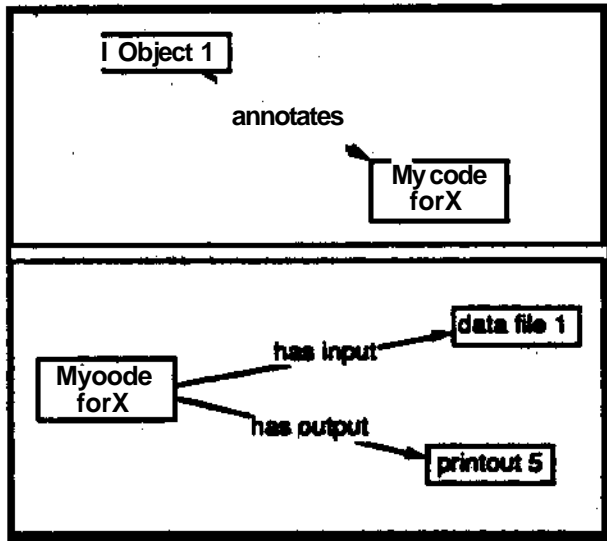


Fig. 3 Nodes and links models typical of those created in n-dim. These models relate to two of the views shown in Fig. 2

A user can also include any number of user-defined, labeled links to express relationships between pairs of model parts (including self links). For example, one could use a link labeled "is-preceded-by" to relate two activities which are part of an activities model. In usual directory/subdirectory file management schemes, such links are not available nor are they explicitly available in a hypertext.

Conceptually objects in n-dim exist in a flat space. One imposes an organizing structure over these objects when one relates them by placing them into an n-dim model. Each object needs to exist only one time but can be a part of any number of models.

Publishing, revision management, prescriptions, and access control: *Publishing* is the mechanism n-dim uses to make an object persist. n-dim uses publishing to maintain history and to facilitate communication. A proper metaphor is the library. Once one publishes an article, it will always exist and cannot be altered. Any two people claiming to read the same article actually have read different copies which are guaranteed to be the same. We can partition the objects to which n-dim can point into two types: those which the user has given to n-dim to store and manage and those which the user manages outside n-dim, such as files in his/her file system or external programs belonging to the system as a whole.

A user creating an object in n-dim will do so by an editing process. To share a model with others, its creator must *publish* it. n-dim will not allow a

published object to be modified by anyone, not even its creator/owner. The user must first publish all the parts in a model that n-dim manages. If n-dim manages everything the object knows about either directly or indirectly through its parts, then a published object is guaranteed to be irrevocable.

n-dim must not support finality since a published object may not be considered final. The original object will continue to exist along with the revision. Anyone can relate a revised object to an original by creating a separate *revision model* linking the two objects with a revision link connecting them. By treating revision links as special, n-dim can locate revisions rapidly.

Often a person wants to open the latest revision for an object rather than the original. To accommodate this need, we created the notion of a *prescription model*. It has two parts: the original object and an n-dim model representing a person or group of persons. If a user clicks on a prescription model, n-dim does not open it. Rather it locates the latest revision of the original object made by anyone in the group and opens that.

Two benefits accrue. First, one gets the latest revision even for published parts and, second, n-dim keeps all revisions as a history of the object. Two people revising a published document can only copy it and edit the copies, side-stepping the so-called *long term check-out problem* with databases and document management systems. They will have to reconcile any differences later.

A user can control who has *access* to any of the objects s/he publishes. A person can only see an object created by someone else if it is published and if the creator gives that person access to see it. A conflict occurs between access and publishing. Access is an attribute of an object. How can one change the access for a published object? First of all we add the attribute of access to an n-dim model by adding a special part to each model that points to an n-dim model of a person or a group of persons. Then we resolve the publishing issue by requiring the access model to be a prescription model so the pointing is indirect. With this approach n-dim knows who currently has access and everyone who has ever had access and when.

Rules: Users want n-dim to tell them if certain events occur. For example, if someone in a group creates a revision link to an agenda model the user published last week. We introduced *rules* that events known to n-dim trigger. Publishing a rule (an n-dim model) turns it on; revising a rule allows one to alter it or turn it off. n-dim remembers a history of all rules and when they were active because they are n-dim models.

Launching external programs: A user can structure data in an n-dim object type we call *frame*. These data can become the input for an external program such as a simulator. For each type of external program, n-dim

must have a special purpose *wrapper* program to translate data from an n-dim frame into a suitable input file for the program. A completed wrapper will also translate output from the program into an n-dim frame. The external program is an object at which n-dim can point. Thus an icon for it can appear as a part of an n-dim model. Double clicking with the mouse on this icon *opens* the external program which in this case launches it by first translating the data and then triggering its execution.

Storage of objects: Each n-dim object has several attributes such as a unique n-dim identification string, the name of an owner, a time of creation, a time of publishing* and a language used to create it. It also has its contents which, for a model, is a list of has-part and labeled links. Graphical objects contain graphics, etc. Text (for example, a whole article), graphical, audio and animated objects can be very large.

n-dim stores all attributes except the contents for large objects in any of a number of databases throughout the web. Large objects are stored as files or, if the database system will allow it, inside the database. We use the ISIS system [Birman et. al., 1992] to connect workstations over a network. Available both as "freeware" until recently and in a commercial version, ISIS takes care of all the messy coordination problems that occur when sending messages back and forth among networked computers. Each participating workstation has a piece of ISIS code on it that sends and receives messages from the other participating workstations. To find an object, one can ask the local ISIS code to broadcast the request to all the other workstations. It passes responses back to the requester.

By storing all the model attributes, one can construct a query to find a model which may be stored anywhere in the web by its attributes. For example one can ask for all activity models which Richard Smith owns and which he created after January 15, 1994. ISIS broadcasts such a query to databases throughout the n-dim web.

Languages: a mechanism to support inductive learning:

A language is a special n-dim model that any user can construct. Using a model as a language when creating a new model limits the new model to having only the object types and link types listed in the language model. A universal language allows one to construct a model using any type of part and any type of labeled link. Using the universal language allows users to try out new types of models which they may then share with others. For example, suppose a member of a design team sees the need for an activity model to organize his activities. He will create instances of such a model allowing it to contain only atomic activities (as frames containing activity names, start times, etc.), other instances of activity models and precedence links. He may share this with others in the group who adopt it.

One of them may write a method to find critical paths. Finally, to pass model type and its supporting method to others, someone in the group creates a language for constructing a legal activity model. Until the language exists, a model author enforced these restrictions manually. We see the capabilities evolving with time in this manner. The generalization process needed to hypothesize the usefulness of a type of model and then to create a language to pass it to others is an *induction* process, n-dim does not do induction, but it supports its users in this process.

Using information management technology

We shall discuss three key uses for information management technology: to generate standards, to support collaboration, and to permit an organization to reuse and to improve its work processes.

We have argued above that we can use databases, document management systems, the World Wide Web and n-dim for information management. Since n-dim uses databases to store the things it controls, we view it as the software one will add to a database system to effect the type of information management we have described above. We, therefore, shall generally limit our comparisons here to using n-dim vs. using the WWW. We believe the following to be the key comparison issues.

Structuring Information

The WWW is a hypertext and, as such, supports the has-part link but not the notion of a labeled link directly. A user can group items by referring to them from a single page. The text on the page can serve as the annotation for the grouping, describing why each item is there and how all the items relate to each other. Thus one can informally set up the equivalent of semantic nets that underlie the design of data structuring in databases. This capability makes the WWW very useful for structuring information, n-dim, with its labeled links, allows the equivalent structuring to occur.

Allowing Ad Hoc Views

There are two extremes for allowing users to choose how to structure information: complete anarchy or complete control. Anarchy allows users to discover new ways to structure the information - i.e., it supports induction, but it precludes users easily attaching analysis codes and using the data as input. Complete control supports attaching of analysis codes, but it restricts use to routine, predefined structuring of information.

The WWW is complete anarchy. Users are free to structure information anyway they wish. At the other extreme is the complete control database managers often impose on users of corporate-wide databases. Typically, experts have predefined all the schemata allowable for users to enter data. For some database managers, users

can invent their own schemata for their own data in their own space but cannot readily share it in that form with others, n-dim supports both approaches. Languages other than the Universal Language support predefined structuring while the Universal Language allows complete anarchy. Users can share all models. Schemata, like other n-dim models, can be redefined.

Finding Information

A major activity in an IM system is locating things which are useful. Searching for interesting things in the World Wide Web has generally been by browsing (called *surfing* by some) around the Web. One can place the names of interesting pages in a "Hotlist" that the interface program Mosaic maintains for each user, making it easier to relocate interesting objects.

With the development of a new program called the WWW Worm (WWW) [McBryan, 1994], one can now also locate objects using text strings that occur in their title, their URL, or in any button text and URL they contain. The WWW runs in off-peak hours¹. It discovers more and more HTML files by recursively scanning through all HTML files of which it is aware. For each HTML document, the WWW tabulates its title, its URL and each URL the document mentions together with the text annotating that URL.

n-dim stores several attributes in searchable databases for all objects it maintains. In particular, it keeps all the links contained in a model as part of its searchable attributes. Thus one can search over all has-part and labeled links, owner's name, modeling language, date of creation/publishing, and so forth, to find things of interest. Also one can simply browse.

Maintaining History

The WWW has no specific mechanisms for maintaining a history of the objects in it. Participants can edit any object they own, and the new version automatically replaces the old provided the author alters neither its name nor its location. Thus one will always access the latest version of an object. A group of collaborating users of the WWW can agree informally to maintain a history by the following mechanism. Before editing a document, copy it into a file with a name such as fileX.rcvision12. Then point at this revision in the new version. This form of revision management is limited to a chain of revisions that the owner maintains.

Many databases also do not have a history mechanism, though some newer ones now keep everything that has ever been put in them, and users can

ask for the answer that would have been given at noon three days ago.

A main feature of n-dim is that published objects are both immutable and permanent. In principle* no one can change them nor remove them (of course* a superscr or loss by failing to backup information can defeat any such mechanism), /i-dim users must publish objects to share them so they will publish lots of things. n-dim has a complete revision management/recording capability. Anyone can revise any other model, even those they do not own. Prescription models allow one to open only the latest revision for an object* allowing one to find only those revisions made by an identified group of people. Access control uses prescriptions so n-dim records who has ever had access and when.

Speed of Access

The WWW and most databases have each object stored in one location* Times speed of access depends on how rapidly one can copy the object from that location. The publishing paradigm of n-dim guarantees the immutability of objects. Thus it can keep multiple copies dispersed strategically around the web, offering the opportunity to speed up searches. For example, n-dim can place copies of objects one accesses frequently in one's local database.

Controlling Access

The WWW defers access control to that of the operating system for the computers on which the information resides, n-dim maintains its own access management system on top of that for the operating system. We described it previously.

Controlling Who Can Join

One of the beautiful characteristics of the World Wide Web is that no one is in control of it nor who can join it. Anyone can join at any time by fetching and installing the right software. New users must make their files accessible and both accessible and known if they want to be a source of information for the community, something which happens when others point at their objects.

Through a policy decision, n-dim, like a database community, controls who is in and who is not.

An Application: Generating Standards

Standards generation is a process by which the practice and experience with a product designed and used over time is codified into a set of rules, representations and methods. Information management technology provides a means to aid this process by allowing the creation and co-existence of alternate models of standards in the same

¹ To find the WWW, open the URL <http://www.cs.colorado.edu/homes/mcbryan/public.html/bb/summary.html> which is the home page for the Mother of all Bulletin Boards.

computational environment. In a product design organization, this technology creates a new possibility where standards creation is not an ad hoc periodic process but a continuous process where the practitioners can incorporate changes in the social and technical requirements as clearly and as fast as possible to minimize variations in the designs produced.

The WWW is not intended to support multiple classifications and the history that are critical to the standards generation process. However, a standards group could get input from a large external audience by making the standards documents available for comment as we do today on paper.

An Application: Supporting Collaboration

Imagine a group of designers working together. They generally work at their own desks, meeting together every two to three weeks. The design is something new for them. Their first tasks are to gather information, establish a working vocabulary, and define and create a common understanding of the problem and possible approaches to solving it. They agree to scan images into the computer of all hand written documents and to save electronic objects such as e-mail messages, drawings, and so forth. Databases, documentation management systems, the WWW and n-dim offer very useful ways to keep and structure this information. By using the protection schemes on their computers they can protect their files from others.

n-dim is our hypothesis for the type of support needed for collaboration. Our next step is to prove this hypothesis by actual case studies. Our beliefs are that maintaining history is important (it is frustrating to have an object you count on disappear from sight, even one that has errors in it), ad hoc models are a necessity as no one yet knows what information they will collect nor how they want to structure it, and being able to find things by attribute is crucial - browsing is not sufficient.

An Application: Reusing and Improving Work Processes

To reuse and/or improve a work process, we believe one must maintain a high fidelity history of instances of that process. The company can then use the history as a starting point for reuse. It can also study these recorded instances to discover what really occurs and to look for ways to make improvements. Interviewing people after they complete an activity does not give an accurate view of what they have done as they will rationalize their actions. One can demonstrate this assertion by video taping a problem solving session and comparing it with a follow-up interview.

In conclusion

Greatly expanded information management is now possible. It will take time to discover all the attributes support systems must have for information management as no one has yet really tested these systems. But we shall. This is an exciting but in some sense troubling time as shown in the next section.

A Small Glimpse of Etopia - an Information Society Some Time Far into the Future

Past attempts to characterize an ideal world include Sir Thomas More's *Utopia* written in the early 16th century, Edward Bellamy's *Looking Backward* and William Morris's *News From Nowhere* in the late 19th century. They offered solutions to perceived abuses in the cultural and social systems of the times. The value of these solutions in these works continue to be debated but many have been adopted; the credit cant from Bellamy and self education from Morris are examples. We thought that it might enliven discussion to speculate about the future of process design in an hypothetical future electronic information age; we call *H Etopia*,

In Etopia work falls into three categories of activities: 1) Education, 2) Business and 3) Consulting. Business activities include the transformation of material or knowledge inputs into outputs. Businesses are privately held reporting to stockholders or publicly held reporting to Oracles. Educators consolidate what is known about reality into what is necessary to teach students about how to apprentice in specific aspects of business or consulting. Consultants provide the service of conjecturing about what is presently unknown to businesses. Businesses provide material needs, social welfare and public infrastructure to all inhabitants. Representatives from each of the three work categories are elected by their underwriters to manage businesses. An agent, called Network, provides all tools and data about all work activities for all institutions in Etopia. Each institution can maintain some data and a number of methods for confidential use (similar to other privately held property). Businesses and individuals pay taxes based on both 1) the amount of data, methods and other capital assets which they hold privately and 2) on the value their business adds. The taxes support Network and the social and infrastructure activities.

There is a capability for all to share information from publicly funded research. However, the Oracles have fixed delays for access according to type of information (such as medical or defense types) as well as to which operational organization the agent seeking the information belongs. At present there is a maximum 18 month time delay. Initially it was four years, but, as the evaluation systems improved to recognize ideas that truly benefited society, the Oracles decided to shorten the delay period.

Network automatically translates all information into different languages including nuances and syntax of meaning. Where meanings are ambiguous, as was often the case in previous technical papers, Network asks the author for clarification.

The Etopian society fosters teamwork by rewarding and recognizing outstanding team performance. All members of the team receive the same level of reward. Contributors learn to work together without concern about individual glory. Dissension among team members is accepted, but Network evaluation processes aid the rapid formation of a consensus.

The society encourages personal mastery of special disciplines. It also promotes Network dialogue among people who are working on similar ideas. Network records who invests first. The Oracles establish criteria for novelty. Oracles are elected from three highest levels of apprentice and full consultants. Network makes the election based on the record of the value it gives to the consultants' judgments. An Oracle's term of office ranges from three to twelve years, based on random assignment at election.

Etopians have solved the problem of ownership of ideas. Network monitors the misuse of or taking credit for others' ideas, and Oracles punish offenders by ostracizing them from the information community. Network, with oversight by the Oracles, patents and values ideas. Institutions and individuals, as users of ideas, remunerate idea owners and their organizations. The society encourages the use of others' ideas.

Education in Etopia is truly a pleasure. For education through the BS level, students and mentors view the same virtual screen. However, students communicate verbally through Network by asking questions about the subject being discussed. Instructors have provided several levels of help in the form of a hypertext. Network allows users to select a depth or a breadth view through the hypertext information tree. As the deeper levels of information become more esoteric, Network automatically supplies required underlying knowledge and theory needed to interpret the given information. As the level deepens further, Network opens actual experimental databases. If the user selects breadth, Network provides associations with other similar branches of knowledge. If depth in a hypertext is analogous to a dictionary, breadth is analogous to a thesaurus. Educators learn because Network gives them instantaneous feedback on the questions students ask, the paths they take to solve problems and on whatever Network perceives is missing. They can interrupt student learning sessions to provide guidance. Educators continually refine the information trees. The breadth dimension of the trees has grown so that all former educational disciplines now overlap. The core curriculum, always the center of much controversy, now consists of a path connecting the nodes of a few fundamental subjects. Students no longer attend lectures

but are connected to a virtual interactive discussion session when they complete the exploration of a node.

Society encourages all citizens of Etopia to take risks with new processes and products without undue concern for liability. Initially, all citizens were concerned about the ability of Network to record timing and ownership of all activities. This led to a reluctance for users to take advantage of the capabilities of the system. Decision makers feared that knowledge about who made which decision would lead to legal liability far into the future. Someone wrote a novel about the family of an engineer who lost everything in a class action suit because Network had recorded that this engineer made an incorrect judgment. He had tested a new ceramic material for a bearing based on superior performance in the laboratory over a three year period. The ceramic actually underwent a very slow decomposition reaction in the presence of some newly introduced lubricant additives. As a consequence of this slow reaction, the wheel bearings on several very high speed passenger trains failed over a period of a week. There was significant damage and loss of life. It turned out that the reaction in question had been mentioned in passing in one of the Network documented courses that the engineer had taken in industrial chemistry seven years before. In the novel the Oracles agreed that the engineer should have known about the reaction. Because of public outcry about the novel, and the recognition that these concerns were real, legal rules were rewritten to limit both the scope and the time that historical data from Network could be used for evidence in litigation as well as the evaluation of personal performance. Furthermore, the capability of Network to associate and evaluate marginal information was enhanced so that it was less likely for information to be obscure. Only rarely now are legal questions about liability tried before the Oracles.

Etopians appreciate change. For eons humans and their evolutionary predecessors were conditioned not to accept new things into either their everyday life or their thinking patterns. A random selection of inhabitants now test new products and ideas. It has become part of the cultural entertainment to be a participant in any testing program.

Etopians understand well the importance of complete system modeling. A Model Building Tool (MBT) reached a high state of development as a result of great attention by many consultants. MBT suggests all potential constructs for the system being modeled. As the user pares the structure, MBT automatically generates a list of assumptions. Thus, the documentation for any model contains the list of underlying assumptions. A surprising result from using MBT is that in some cases model building is an initial valued problem. Some Consultants have found that slightly different models result depending on the urgency of the situation for which the model is needed. This significant finding is the center of much consternation in

the Consultant community and an active area of investigation.

Professional societies have melded together to form a consortium of specialized technical consultants. Multiple hierarchies of consultants each have their own bulletin boards for communication. Consultants are trained by taking virtual courses in their specialty. The experiences in these courses include emergency situations from actual equipment and system operating history. Societies use licensing exams to qualify consultants, extending the BS education system. They score apprentice consultants based on the speed, originality and soundness of their responses and actions. Consultants' rankings are based on the same system as "Go" players. There are 30 levels of entry apprentice, six levels of apprentice and nine levels of full consultants. There are only three or four consultants at level 9 in any field at any one time. Project Managers bid for consultants at any point in time with higher level consultants being more expensive

Network provides immediate costs for equipment, labor, modeling, design, laboratory work, etc. Network establishes consulting and business fees and adjust them according to demand, utility and uniqueness of products. Inter-business supplier-purchaser relationships have been firmly established on the basis of contractual auditing procedures. These procedures allow the purchaser to receive certain information about the cost of manufacturing or otherwise creating products. This allows purchasers to make better decisions. One result is that now each vendor supplies to the purchaser the best estimate about cost of equipment items based on the probability that equipment will actually be purchased, the delivery schedule needed, the costs of labor and raw material and the process and mechanical specifications. Vendors also provide information about costs and characteristics of alternative equipment which they have under development, thus assuring a plant will be built or revamped with state of the art equipment.

Network has been given the capability to act as an independent "third party" to search certain aspects of all institution information bases. Thus Network can be used to automatically "benchmark" the performance of similar institutions, point out the differences, and recommend improvements. Where it finds similar data (for example corrosion data for various materials in various chemical atmospheres), data bases are adjoined, greatly improving the accuracy of models used for equipment design.

These are only a few of the fascinating aspects of Etopia that relate to process design. Indeed, Network is in the process of writing a virtual book about Etopia.

References

Bellamy, Edward, Looking Backward, Edited by Cecelia Tichi, Viking Penguin Inc. New York (1982).

- Birman, K., et. al., The ISIS Distributed Systems Systems Toolkit Programmers Manual ISIS Distributed Systems, Inc., Ithaca, New York, (1992).
- Conklin, J. and Begman M. L., "gIBIS: A Hypertext Tool for Exploratory Policy Discussion," ACM Transactions on Office Information Systems, 6(4), 303-331, (1988).
- Corrigan, James P., "Is ISO 9000 the Path to TQM?" Quality Progress, May (1994).
- Drucker, Peter F. Management-Tasks, Responsibilities, En&icjg, Harper and Row, New York (1973).
- Grønbaek, K., Kyng, M., and Mogensen, P., "XSCW challenges: Cooperative design in engineering Projects," Communications of the ACM 36, 6, 67-77, (1993),
- Hammer, M. J. Champy, Reengineering the Corporation, Harper Collins New York (1993).
- Lampport, F. Latex: Users Guide and Reference Manual, Addison Wesley, Reading, Mass, 1986.
- McBryan, O.A., "GENVL and WAVV: Tools for Taming the Web," in Nierstrasz, O. (ed), Proc. first Intl WWW Conference, CERN, Geneva, May (1994).
- More, Sir Thomas, Uigajk Translated and Edited by Roger M. Adams, W.W. Norton & Co. New York (1975).
- Morris, William, News from Nowhere and Selected Writings and Designs, Edited by Asa Briggs, Penguin Books Ltd. London (1962).
- Naisbitt, J., P. AbuTds**A Re-inventing the Corporation, Warner Books, New York, (1985).
- National Research Council, Improving Engineering Design--Designing for Competitive Advantage, National Academy Press, Washington, D. C. (1991).
- Peters, Thomas J., Robert H. Waterman Jr., In Search of Excellence, HarperCollins, New York (1982).
- Scherkenbach, William W., The Deming Route to Quality and Productivity—Road Maps and Roadblocks, CEEPress Books, Washington, D. C. (1986).
- Schmidt, K., "Modes and Mechanisms of Interaction in Cooperative Work," Edited by Caiia Simone and Kjeld Schmidt, COMIC Project Deliverable 3.1, 21-104, Computing Department, Lancaster University, Lancaster, U.K., (1993)
- Senge, Peter M., The Fifth Discipline: The Art and Practice of the Learning Organization, Doubleday, New York, (1990).
- Taylor, Frederick Winslow, The Principles of Scientific Management, Harper, New York, (1911).
- Thayer, Ann M., "Chemical Companies See Beneficial Results from ISO 9000 Registration", C&EN, April 25 (1994).
- Unilogic, Scribe Document Production Software: User Manual, Unilogic Ltd., Pa, (1985).
- Walton, Mary, The Denting Management Method, Dodd, Mead & Co., New York, (1986).