

ATARI
Systems Research

----- INTERNAL DRAFT ----- INTERNAL DRAFT -----

To: Alan Kay, ASR *JK*
From: Stephen Weyer, Alan Borning, Dave McDonald, ...
Subject: Encyclopedic Research Plans
Date: 1 July 1983
Filed: [weyer]PROPOSAL.MSS, .LPT

Comments, please.

Steve

1. Preamble

Our intention is to identify key issues and questions and to suggest leveraging projects that hopefully anticipate or feed into other projects and overlap several research areas. The "project" we have in mind is not a monolithic one but rather a confederacy of projects related by common themes and integrated eventually into a system. Although the final picture may not be totally clear, we believe that there are initial directions that we can pursue that will give us needed experience as well as useful research results.

This is written for ourselves and is intended to allow us to get started right away. Because it assumes familiarity with our shared context from earlier discussions and memos, it should not be taken in isolation as the project description or a set of justifications for management. In addition, we are more concerned with the partial ordering of many related tasks rather than absolute project completion dates for any one of them.

2. Metaphors

We have attempted to understand our shared goal of providing rich environments for learning, exploration and entertainment by looking at examples of these individual environments, our imagined electronic realizations of them and the experiences of individuals in them. Metaphors or themes such as sharing and fantasy, communication, and learning through exploration (as well as education) have been used to explain and generalize across environments such as intelligent encyclopedia, information community, knowledge spaces, media living room, amusement park, art and science museum, interactive movie, and adventure game. It is possible to argue interminably about which theme to adopt as our exclusive banner, or which environment might be an

"umbrella" project to subsume the others. Far more valuable would be to postulate (like the nine blind men perceiving an elephant) that what we are proposing has important aspects of all of these -- it may be a hybrid creation that neither closely resembles nor actually replaces any of its traditional analogs (like books or amusement parks) but provides some new and synergistic opportunities.

We should be selecting projects that express more than one theme but avoid worrying prematurely about how to do everything at once. Rather than building "just an encyclopedia," for example, we should be thinking instead about building systems that embody "encyclopedic knowledge" or convey an "encyclopedic perspective."

3. Initial Project and Tasks

The system we envision will provide opportunities for representing, browsing, and otherwise experiencing knowledge in the form of simulations and information worlds (databases) and presenting graphical and natural language explanations of the knowledge -- representing, accessing, and presenting information. Although highly flavored by the encyclopedia theme (in its traditional and electronic forms), there are many other aspects that we may want to focus on as well. At the end of this memo appears a list of phrases mentioned or detailed in earlier meetings and memos that can provide us with reminders about issues that we might want to consider in specifying and creating an encyclopedic system. By themselves, the phrases are not detailed issues or research questions though they are intended to suggest them.

3.1. Browsing in the Encyclopedia Database and Knowledge Environment

One initial approach is to take an existing encyclopedia database such as Grolier's Academic American Encyclopedia (most likely), or Encyclopaedia Britannica's Encyclopaedia Britannica or Compton's Encyclopedia and to enrich and enliven it along a number of dimensions. We propose to focus on the creation of browsing and searching mechanisms that will initially resemble Findit, Rabbit, hypertext or inspector-type window displays, and that will be connected to an existing database (backend) system such as Ingres.

With respect to specific encyclopedias, it's likely in another month or so that we can obtain tapes from Grolier's for Academic

American (AAE) text articles, index (pointing to pages, unfortunately) and tables. If we can somehow keep our options open with Encyclopedia Britannica (say, by arranging that the Grolier agreement is through Warner's TIES project, with ASR consulting on interface research), then we might be able to obtain Compton's and EB later on.

Although an encyclopedia will be a primary knowledge domain, we intend to apply and extend our approach in other knowledge-rich areas as well, as other opportunities arise, such as the Lisp Machine environment (flavor browsing and online access to manuals, documentation and examples), general work environment (personal notebooks, calendars, documents), foreign databases (e.g., Dialog, Nexis), library catalog and sources, and other databases such as the Whole Earth Catalog and National Geographic pictures.

The system will not be "just" an encyclopedia, in either content or style of access. The human "readers" are really "searchers" and "writers" who are using their personal notebook to jot down questions, observations, write homework assignments and papers, share results and messages, etc. Behind or linked to the notebook are the encyclopedia and other information resources we make available. The encyclopedia is not cast in stone, not because it is electronic and can be updated, but because readers can augment their own copies by adding their own marginal annotations, bookmarks and crossreferences and by connecting to other databases such as a dictionary, current world events and online discussion groups and journals. The system should be able not only to augment the content by knowing something more about the subject than is obvious from the text alone (the realm of AI), but also to aid the search process by commenting to the reader on where and how to look.

We would like to find out how using a subject index and thesaurus compares to a free-text access approach. How can we follow the many paths through this highly interconnected knowledge without losing track of where we've been and where yet we'd like to go? These are questions that require experimentation and presuppose the facility to collect protocol data on user actions for later analysis or replay. How can yet other knowledge sources (e.g., a dictionary or other persons' comments) be brought to bear on the question at hand but without distraction? How can information be shared -- not only packaged information but also real-time information as several people explore an information environment together? We want to integrate photographic images and computer-animated dynamic pictures into both the content of articles and into the structure of the system itself, e.g., as browsing maps. We want to provide

because of the many possible ways they could be viewed and transformed: scale, 2-d projections, kinds of features (mountains, mines, tribes), boundaries (geological, political), user-defined filters, etc. There is some information in the encyclopedia index about latitude and longitude of locations. About obtaining and drawing the maps themselves, one initial possibility is to use some scanned images; another is to obtain map databases from elsewhere (the commonly available CIA database of world coordinates is probably not adequate).

3.3. Using Simulations

One of the things contained in encyclopedia articles will be simulations. For example, in an article on the solar system, one might find an orbital mechanics simulation; in an article on mechanics, a resonance demonstration; in an article on economics, some simulations of different economic models; in an article on geometry, an animated demonstration of a geometric construction.

We have constructed a few representative physics simulations of such things as resonance, colliding particles, and orbital mechanics. These currently run in Interlisp-D on a Xerox Dolphin at the University of Washington. We are in the process of porting these to the LM-2, so that they can be integrated with a prototype browsing and retrieval mechanism and can be presented along with appropriate text.

In the longer term, however, to make the simulation facilities useful, it must be convenient for authors to construct, test, and modify their simulations. The authors will usually be experts in their own fields, not in programming, so the environment for constructing simulations should not just consist of Lisp or some other language. Further, we should make it possible for readers of articles to build their own simulations by making use of building blocks provided with the articles, so that in addition to its other qualities the encyclopedia becomes a kind of laboratory.

To accomplish this, we are starting with the ideas in the ThingLab simulation laboratory. Examples of simulations constructed in ThingLab include electrical circuits, mechanical linkages, bridges under load, and animated diagrams to demonstrate geometric theorems. A basic notion in ThingLab is that of constraints. A constraint specifies a relation that must be maintained. Keeping that relation maintained is left up to the underlying system, rather than being the responsibility of

the user. Constraints are thus useful tools for constructing systems for use by people who aren't computer sophisticates, such as the prospective authors of our encyclopedia articles.

A successor system to ThingLab is currently being designed. Differences from ThingLab include:

- provision for constraints describing dynamic systems, e.g. an orbital mechanics demonstration
- provision for users to construct new kinds of constraints in a convenient, graphical way (since we won't be able to anticipate all the building blocks required)
- a strong notion of multiple views on an object or simulation, with ways of constructing and composing filters to generate these views.

Implementation of this successor system will be in two stages. The first step will be the implementation of a prototype system in Smalltalk-80, running on a Dolphin. This implementation will serve to make these ideas concrete, but will be done with little concern for efficiency. It should, however, exhibit the sorts of interactive graphics capabilities that we want, so that we can use it to generate pictures for papers (and perhaps to make a videotape by suitably varying the recording and playback rates). Programming on this version will begin this summer at the University of Washington, with the work being funded from several sources; the system will be available for ASR to use.

Based on the experience gained with this implementation, a second system will be constructed. The host machine and language will be chosen for compatibility with the other parts of Atari's encyclopedia system; if Symbolics 3600's are used, it will be implemented in Lisp on that system. This implementation will be written with efficiency as one of its goals. Further, the viewing and filtering techniques will be integrated with those used in the rest of the encyclopedia system.

3.4. Natural Language Explanation

Articles in the present encyclopedia are "fixed" on a single presentation -- the text that their author originally wrote for them. We want (eventually) to be able to back up from the articles to a purely concept-based knowledge network that is flexible about how it is presented, tuning its choices of

language, pictures, simulations, film footage, etc. to match the needs and abilities of the audience it is working with. Achieving this requires advancing the state of the art along two dimensions: knowledge representation to get the information into the knowledge base in the first place, and interactive presentation to get it out in a deliberately tuned way. There is certainly more than 10 years' worth of work required to accomplish this, so we need to set some near-term goals and to identify the kinds of basic research that will be required.

One immediate "probe" we expect to do is a careful "reverse engineering" of a number of encyclopedia articles. This will serve two purposes: (1) it will tell us if the present art of knowledge representation is even in the right ballpark for the encyclopedia, and (2) it should move us significantly along the path to an "intelligent Propaedia" by identifying the kinds of "organizational glue" that hold together short encyclopedia articles and paragraphs, potentially allowing us to tear them apart and refashion them as (somewhat) tailored presentations even before we can properly represent the root content information they contain.

The mid-range problem (months to a year) is to get a handle on how to manage an explanation that can be carried on at different levels of conceptualization. Critically, the explanation must be interactive: users get to point out what they don't understand and where they want to know more, and it must be multi-modal: a deliberately selected, dynamically created combination of voice, diagrams, active graphics, background visuals -- whatever seems to best convey the information in the most pedagogically effective way. Beyond that the choice of subject matter is completely up for grabs: almost any domain will present interesting problems, and the best domain will be whatever one has the most people caught up in it and contributing to it.

3.5. Other Related Projects

"interfaces", Library, tv aquarium, videodisc, music education, dynapad....

3.6. Products

At the outset, we are enthusiastic about the prospect of our research prototypes generating and influencing products (and will be disappointed if Atari resists moving strongly into areas that are strongly communications and learning oriented) -- for example, a decent online encyclopedia (not the "intelligent" encyclopedia) should be possible within 3 years.

We would hope, for example, that we could create a system whose architecture would not cause successive versions of products and protocols to be greatly incompatible or obsolete and that could be used directly to simulate important properties of then currently available display and delivery systems (somewhat analogous to the Games Emulation System). Although we are committed to consulting and consorting with product groups, we do not have the time, resources or experience to pursue actual product development ourselves.

I. A List of Ingredients, Issues and Ideas for Projects

As mentioned at the beginning of the last section, the following is a collection of brief reminders and suggestions. An exercise for the reader might be to add more entries to the list, merge in Alan's recent list of "foci and trajectories", impose a coherent organization on it, and select a likely viable subset for a project.

exhibits/models
kits
simulations, constraint languages
electronic mail, calendars&scheduling
phone books, bulletin board, libraries
personal notebooks, annotations
electronic geodex
whole earth catalog
community memory (project by that name still alive)
encyclopedia (Britannica, Groliers)
pictures (National Geographic)
indexes (e.g., Propedia)
maps/atlas
dictionaries, thesauri
structured documents: headings, tables, crossreferences
online books and manuals (e.g., Lisp, auto transmission)
AI, frames and knowledge representation, Eurisko
programming structures (e.g., flavorarchies)
online library (catalog)
foreign databases (Dialog, Nexis, Kurzweil&OCR)
natural language understanding, paraphrasing
stories, adventure games and movies
humor, joke generator

guides/agents
personality, style
authors, bias, editorial control, censorship
what is brought to system by reader/user
foreign agents for db (KGB = Knowledge Guided Browsing?)
dumb guides, i.e., reactive windows

filters, views, perspectives
reasoning and presentation based on example and analogy
natural language generation and explanation
formal theories of interface and human interaction

tours
search history
annotation, reminders
query by example (Rabbit)
browsing paths, browsing by analogy
tutoring, pedagogy
models of user understanding and misconceptions