

HARCOURT CONFIDENTIAL

Theodor H. Nelson

THE XANADU SYSTEM: A Discussion Paper

23 December 1966

INTRODUCTION

This paper is intended to communicate with both management and programmers. On neither side is the system's practicability assured. Everything described here is possible and, I think, desirable. Whether it is feasible, that is, possible for a reasonable price, is not yet clear. How much of the system described can be implemented on contemplated equipment is hard to tell. Matters of programming cost and time inconvenience to the user may prove to be restricting.

These are some remarks about the XANADU system, a proposed setup for handling a large number of corporate and creative functions on a computer with display screens and a rich file handling capacity. All that is said here will be tentative thought and ideas on how such a thing should work. The XANADU system is supposed to be a general-purpose, highly flexible information handling system for the entire firm, especially suited to experimenting and easy usage. This paper presents design remarks on this master program, and a ticklish job this designing is, too, since it is the executive and editorial staff at Harcourt, Brace & World who must live with it and like it.

This is hard work. It is hard, not because there is a certain exact way to do everything, but because-- to be candid-- it can be done in any number of different ways, and the many options are constrained only slightly by technical bounds. We could choose to pretend that "the computer" requires a certain mode of operation, even a particular posture, of the editor who will use it; but this would be a falsehood. With certain exceptions, it will do practically any darned thing an editor might want, if we are prepared to devote, sacrifice or subordinate this or that feature of the machine to the darned thing we want done.

Very well. The system we want to put up will be sophisticated (its fancy features will appeal to the most arcane); general (by making little arrangements this way and that you can make it do service in many different guises and stances; useful, so that its features are relevant to the comparisons and vantages you need it for; and pleasant, aesthetically attractive, simple, fleet, and fun.

The basic idea is to build one system so thorough and general that every text-handling application we can now think of will operate through it. This means that its underlying structure will be all the more anomalous, and unlike any other systems in this area.

Although it is being designed around a very small machine, it is intended to be one of the most powerful text-handling facilities in the world. The general idea is this: the system is meant to be typewriter, filing cabinet, and manuscript, and simulated printed page, all at once. As such, it will not only make the editorial updating of ordinary texts much quicker, but it must also provide information facilities crafted according to particular kinds of needs.

On this unusual facility any conceivable workspace, file discipline, or information structure may be created, and modified at will; no fixed ways of working will exist, except as available options. The set of text and file commands will not be an alphanumeric language unless desired by the user. The system will thus be of use for any mind-aiding process that uses text, symbols and connections.

The basic idea is to build an information-handling system to hold, display, modify, rummage through, and rearrange text and ideas. Unlike any other system that ever has been, this one's intent is to meet the user more than halfway, scrupulously avoiding the interposing distractions of user-burdening

3

computer languages, special characters, symbolic order codes and the like. Moreover, it will be reconfigurable for different purposes: the user may define to the system the types and interconnection types of his information, the way he wants it displayed, the actions of the system, and the way they are to be chosen. Typically he may define choicepoints as "virtual pushbuttons" on the screen, and have them replaced by real pushbuttons, for faster work, when he is sure how he wants to work.

A unique part of the design of this system will be the assumption that our information does not typically come in strings, as is supposed, but is deeply and widely interconnected. Provision for every kind of interconnection must therefore exist, and "ordinary" text will be treated as a trivial case. Thus the system will provide for the creation of annotations, spinoff versions, linked alternative versions, and general hypertexts. (Hard copy will be swiftly available from a line printer.)

Nothing will be restricted in length. Input strings, sentences, chapters, and even titles may go on and on and on. This philosophy of non-restriction extends throughout the system in various ways.

As presently conceived, this system should go far beyond any existing on-line text systems in its flexibility and power. Editorial usage should be only one possible application; the intent in the design phase is to work out a generality that can be turned to many different purposes.

With all this in mind, I ask you to consider the following setup. These features will not satisfy any use in particular; rather, each user may adjust the settings for his own application, and thus we may configure it for many particular uses.

The present design is general-purpose-- indeed, its applicability is so general that from its raw structure it might not be evident that this is to

9

have anything to do with publishing at all, let alone its initial assignments. That is because it must be designed in primitive forms that can be tailored to any specific needs, and these of necessity need to be so general that any tie to specific applications can be established later.

The screen may be thought of as a telescope or the viewfinder of a camera, and we can zoom it to any place in a text, or text complex, that we want to look. But the importance of the device is in its ability to bring swiftly together materials that need to be put together, or seen together. Not just what is sitting on the screen at a given moment, but also the things that are lurking in the wings, standing by or momentarily displaced, constitute the user's workspace.

The system as presently conceived will allow the user to define the types of information he will deal with, and how it will lurk, move and reply to his actions. It is as though he might design his own special sheet of paper, with little doors, fold-outs and overlays.

It should not be supposed by the reader that this description is a statement of what computers "can" and "cannot" do. Rather, this is the design of the system which may be implemented on a computer, and we intend to circumvent the computer's traditional modes of operation, rather than defer to them.

As presently conceived, the system would be simple to operate, easy to learn, and easily taught and adapted to any person and purpose in the company that is given it to use.

~~From the point of view of the system's actual design, however, these particular subjects are~~

5

Four kinds of work have provisionally been assigned to this system.

Ordinary editing. Little need be said about this; it seems likely that employing the console, considerable time can be saved in the preparation of ordinary books for printing.

Test editing; test data bank. For editing tests, both in terms of revising and rearranging test items and accessing information about the file, Xanadu will probably be more useful than any paper filing system. The same items on different versions of the same test, statistical data on their use, could all be filed in ways permitting swift access to corresponding units. Moreover, since the Xanadu structure would permit extensive cross-indexing and categorization, it could be used for the creation and maintenance of a test item file, whose parts could be selected from and rearranged with ease according to all sorts of different criteria.

Print-and-Bind Orders. One of the more difficult managerial problems in publishing is that governing orders to the printers and bindery: relevant considerations include predicting demand for the finished books, the apparent probability of demands exceeding the print order, the storage costs of incomplete material, and finally the schedules of the printers and binderies to be patronized. ~~All these different forms of information could be maintained in~~ Information on all these different circumstances could be made available for thoughtful perusal for the decision maker; they would be arranged in a Workspace that connected each type of information to the things it bore on.

Simulation of 1500 . There is some indication that the IBM 1500 Instructional System may become an influential system around the country in the next five years. While it may not be worth our while to obtain such a system now, it may be worth our while nonetheless to try preparing materials for it. Under Xanadu, it will be possible to simulate the IBM 1500, and create materials for it, without having

one.

Among the other possible applications of this system are the editing of books, the design of highly complex non-book materials for publication, the facilitation of collaborative work between author and editor, experimental typesetting and making of animated movies, ^{AND} ~~and~~ aiding management in comprehending complex banks of data. Perhaps most importantly, the system should help shorten the time required to get critical jobs to press.

From the point of view of the system's actual design, however, these particular subjects and facts are unimportant. The parts that have to be designed into the system are more like the following:

1. Text editing. Contained texts may be modified swiftly with such instructions as "insert," "delete," "move," "copy."
2. User-defined file structure. The user may define his own information types and links, and the way the machine will move and handle them. These may range from an outline form to multi-part linked units.
3. User-defined workspaces and file disciplines. The user may define such abstract working arrangements as push-down stacks, temporary ~~holding~~ holding spaces, "folders," indexes, and windows into his file. He may even specify overall file disciplines that the system will thereafter hold him to.
4. User-defined screen formats. The user may "carve" the screen into separate areas where his files and workspaces will appear.
5. User-defined branching actions. The user may arrange interchange between various presentations by drawing branch graphs, and specifying the conditions under which the machine is to branch; these may include throttle, pushbuttons or screen pointing. This branch-defining technique may be used either for creating work-systems or for defining hypertexts.
6. User-specified ~~presentation~~ presentation modes. The user may have text segments or streams presented as frames, drum rolls, and streamers; and specify which are to blink, how brightly, and when.
7. Continual Modifiability. The system a person is using may be modified ad lib.
8. Reversible processes. Both text and modification commands may be stored in such a way that changes can be undone as far back as records are kept.
9. Evolutionary text generation. Alternative versions may be spun off, and linked to each other so that their corresponding parts may be compared.
10. Complex indexing. The system will have a complex indexing ability for either straight text complexes, permitting the creation of any number of indexes for a given body of information, from which the user may immediately jump to the information being indexed. This indexing facility permits the use of lists or texts to index one another, so that the user may proceed to a corresponding part in the other text or listing. This may be regarded as "backwards" and "forwards" indexing, but that implies a priority which may not

necessarily exist in the material.

11. Alternative drafts. The system will also have the ability to store alternative drafts of either straight text or text complexes.
12. Search. There will also be a word search facility, with which the user may look for the occurrences of specific words and phrases, tally them, index them, process them, etc.
13. Complex Connective Structures. Such connective facilities as super-footnoting, the ability to attach text chunks to an ordinary text and have them available under many different categories; parallel texts, where ordinary texts are linked side by side so that they may be compared or their relations noted; array formats, where other types of information may be stored in matrices and graph structures; and a cross-jumping facility, enabling the user to jump from a given part of the text to any linked part of the text.
14. Arrange and Printout. The system will be able to sort and arrange the materials it contains, and produce printed reports on demand, even interactively with a user. This will make possible catalogues concordances, quick indexes and annotated manuscripts with some swiftness.

9

The Xanadu facility is meant to fill all these technical needs, and yet fill human needs as well. It is intended to be operated on an open shop basis for whatever purposes the company selects. "Open shop" means that the user will handle and direct the machine personally. This means that its elementary operation must be learnable in a very short time (a few minutes), and that it must be appealing to users. Indeed, the system must go out of its way to be appealing to users. It is safe to assume that persons of literary or editorial background will be among the most cantankerous customers ever to deal with a computer-based information system. In the confrontation of computer and litterateur it is the machine that must bow. All traditions of computer usage will be swept aside, in deference to the user's sensitivity about punctuation, capital letters, type fonts, and his presumable distaste for gratuitous numbers, and the arbitrary assignment of ordinary letters to strange new meanings. Our new conventions will be tasteful even to bookmen; but even these of communication between the user and the system will be subject to revision or redesign, in order to please the user or speed them up to his desires.

By the same token, the system's operation, physical plant, and surroundings must be designed for the user's comfort, and the facility should go out of its way to create communication facilities desired by particular users-- for instance, pushbuttons labelled, arranged and illuminated in a specific way because a user says he needs them, or a little "beep" responding to his actions.

Customized hardware may be fitted into a standardized console. That is, the overall response characteristics of the console are given,

10

but users may request customized hardware such as special pushbuttons or pointing devices to readily communicate with a standardized interface.

While the system will be unlike any in existence today, it has many relatives. It is expected to have many of the good features of Engelbart's AHI system, SDC's SURF system, and Licklider's SYMBIONT.

It subsumes INFOL, the information retrieval language developed by Olle at the Control Data Corporation, which is particularly concerned with the categorization and ordering of informational chunks, and their coding, as well as DOL, the Document Oriented Language, proposed by Mandalay Grems of Univac.

In addition, it subsumes in general function a large number of the kinds of work with texts now being conducted in scholarly settings. As a fair indication of this, we may consider the computer programs announced as available in the second issue of Computers and the Humanities. (Of the 51 programs listed, about 27 have their general functions included in this system.)

The system as presently designed will have one feature not currently available in any information system anywhere: the provision for the spin-off of alternative arrangements and drafts of anything; for their gradual shaping and the postponement of final deciding between them.

XANADU FOR BEGINNERS

At the beginning level, the Xanadu system is first encountered as a ~~looking into~~ screen upon which any part of the text may be shown, in both upper and lower case and with some quality of image, and where the user may perform editorial operations upon the text swiftly and with ease. These operations correspond to the conventional operations employed by a proofreader, such as the deletion of material, the insertion of new material (which may be typed in on a keyboard), the moving of sections around the page (ordinarily signaled by a long horizontal curlicue), and the division of the material into lines and pages.

In other words, at the beginning level, the user will have hardly anything at all to learn. The standard editorial operation, or whatever operations exist within a given workspace, will be easy to learn. These functions will be initiated by manipulating pushbuttons or "virtual push buttons" (glowing spots on the screen with labels as to their meanings). The beginning^{er} on the Xanadu system will also have the opportunity of working through these particular work spaces and file disciplines. The basic editorial operations can be learned in minutes; the functioning of a given workspace in perhaps half an hour.

ADVANCED XANADU

While Xanadu at the beginner's level is rather like an extremely fancy piece of paper, advanced Xanadu permits the user to create, as it were, vast information structures hanging in space, connected any which way. The arrangements and facilities available at the advanced level (tentatively called "the Works"), will enable the user to determine the availability

and manipulations upon information, including texts. The advanced user may define information types and create branching structures of response.

The advanced user deals with items, relation and clusters. He may create new entities composed of these objects, and he may direct the machine to create new functions, ways the machine will automatically operate upon data stored in the various formats described. He may define them generally, in terms of types of information and types of choices points, or specifically, creating a hypertext.

~~The advanced user deals with items, relations and~~

At this level the user will be able either to create complex or compound text structures (hypertexts), or himself to create ~~complex or compound text structures (hypertexts)~~; ~~or~~ workspaces for particular uses, arranging choices and carving screen formats arbitrarily according to his own designs.

By "workspace" is meant some arrangement or division of the screen, some particular way of having given informational materials stand by, some particular ways of arranging information of a given kind that will be taken as standard for a given purpose, and whose consistent use the system will enforce. He may create arrangements of particular information within these information types, choice points (and specifications of how the choice is to be made -- whether by push buttons, virtual push buttons, or other means), arbitrary screen formats ("carving" the screen into separate sectors, not necessarily rectilinear), and functions -- collections of operations upon information and information types made available within the Works, which he wants to accrete into compound actions that may be selected or called arbitrarily.

This advanced language, too, will be made no harder to use than necessary. Although the machinery underneath which is to implement these functions may be technically called a computing "language," in the user's encounter with it there is little vocabulary or syntax to learn; as much as possible is relegated to screen pointing, virtual push buttons, and discursive explanations