XANADU SPACE, 1993

Theodor Holm Nelson, Founding Designer, Project Xanadu
and Distinguished Fellow, Autodesk Inc.

Autodesk, Inc., 2320 Marinship Way, Sausalito CA 94965
Telephone 415/ 332-2344 x4679 Fax 415/ 289-4773

THE ORIGINAL IDEA

In the fall of 1960 it seemed to the author that computer screens would become the center of all human information work, whether creative or business-oriented. All informational objects would migrate into digital form and thus become available for use in nonsequential, nonlinear forms. An ever-increasing proportion of documents and works, eventually all, would become branching and thus interactive. Branching and interactive materials would become the world's new generalized form of literature-- for art, science, scholarship, teaching, news, analysis, business reports, science and general presentation of information. This would require many new forms, styles and arts of interaction and presentation.

This would mean in turn that a global-scale digital repository and feeder network would have to be distributed throughout the world. Annotatable, shared material would then be rapidly available to all users, speeding up the flow of ideas and improving their exchange: world-wide groupware, in a way, for attacking all problems. But despite some predictions by enthusiasts of artificial intelligence, all this work would still be guided by human creativity; and this meant human motivations would require the continuation of many familiar human methods (including copyright).

Finally, in particular, we would need a way to track the multiple re-use of the same material in different places, both private and published. In private usage this would help us keep track of the interrelations within our correspondence, projects and other work. In a digital publishing system it would clean up issues of copyright, since material from one document could be virtually included (transcluded) in another through the use of special pointers, with proportional royalty sent to the original publisher whenever quoted material was sent to a user (1).

From the beginning this was conceived not as technology but as *literature*: a fundamental enabling system for the necessary literature of the future, stemming from the fundamental literary paradigm: *all information has a point of view.*

SOCIOLOGY AND COMMERCE

Design of such a system goes beyond mere software. A digital publishing system must take into account the different roles and parties involved: the parties participating, the parties affected, their systems of motivation. We assumed that the parties to large-scale electronic publishing would be the same ones as always: authors, publishers, "printers" (now to be providers of the generalized storage and delivery service). In addition, annoyed and offended adversarial parties, whether private or governmental, will undoubtedly appear, and their concerns will affect the system's world environment.

We also believe that the transclusion feature-- the ability to mix parts of already-published documents with great flexibility and without negotiation-- will have tremendous power in all areas of human endeavor, enabling different parties to create better explanations, presentations, documentary films, music and much more. Thus the transclusion feature will probably create a broad general good. But there will be parties who may not like the way their work is transcluded; for this reason what they are agreeing to must be clear to all participants beforehand.

TOMORROW'S ISSUES, DANGERS AND WARS

The dark side of the new information world probably will stretch on without limit, as has the dark side of the information world we have already known. While many people say that "information should be free," this means two different things": "without cost" and "without restriction."

Freedom 1. "Without cost." The fact is that information has always been expensive, and its possession has given decisive advantage to participants in every field of human endeavor. Plainly there will always be costs associated with information; nothing can make it without cost. The digital publishing scheme that we envision will provide a means for the *easy selling and buying* of parts of works electronically.

Freedom 2: "unrestricted." But the other sense of "free," meaning "without restriction," has important consequences. Tomorrow's issues will have a great deal to do with restricting, and not restricting, information. All the issues of freedom of the press, and of privacy, will surely reappear in the worst of their old forms: censorship, book-burning, spying and other forms of tyranny. A new era of permanent information warfare-- not just among individuals and companies, but between governments and their citizens-- will surely develop, with fraudulent evidence of all kinds fabricated on all sides.

Document fraud. In the past, we could tell the difference between genuine and false documents with relative ease. That is about to change, because the kinds of circumstantial evidence that told us a contract was valid, or a letter had really been written-- the characteristics of paper and its aging, for instance-- will be lacking. Not only text, but now photographs and soon movies, can be counterfeited with less and less difficulty. No objects will be intrinsically trustworthy, any may be fraudulent. This means that techniques for data authentication and encryption will soon acquire great importance for our personal and political lives.

Balkanization. Perhaps the least obvious is the problem of *Balkanization*. This term has traditionally referred to the breaking up of a society-- a system of communication and commerce and sharing-- into destructive factions and rivalries. We may apply the term "Balkanization" to the world of electronic documents as well. Our coming world must have not merely compatibility and connectivity in the usual electronic senses, but *compatibility and connectivity in the new forms of*

documentary interconnection. A world of documentary Balkanization will be far inferior and far less pleasant than what we hope for. I believe that if the world of electronic documents is not unified, the chance of human survival to the twenty-second century will be greatly lessened.

THE ACTION PLAN

Accordingly, the author could think of no more useful work than to design server software for the delivery of hypertext and hypermedia and compound documents to interactive computer screens, in fine-grain demand softcopy form (selling individual fragments on demand) with fine-grain transclusive capability (transclusion among published documents down to the level of selected bytes), upgradable to a large-scale publishing version intended to support open hypertext publishing and open hypermedia publishing with royalty.* A large part of the motivation was to do this right and avoid the kind of Balkanized, specialized and restrictive delivery networks that might otherwise come about.

This work has continued for thirty-two years.

THE SERVER

Thus the Xanadu™ Hypermedia Server was conceived by the author in late 1960 as a system for the digital management and delivery of interconnected documents in all media. In particular, it was to maintain links, multiple versions, alternative versions, and transclusions (logical inclusion without copying) of materials among documents and versions. Originally it was designed for only a single user, with the intent of its being later redesigned for a network. However, successive designs with various collaborators have led to software designs incorporating more and more aspects of the intended megascale publication server. The design now being completed is missing only the back-end internode protocol and the royalty feature.

A number of constraints have at all times guided the server software design:

---

* Note, however, that this overall idea has been quite difficult to express over the years, particularly since so many necessary terms did not exist when these ideas began, such as "interactive," "groupware," "server." Other necessary terms, notably "hypertext," "hypermedia," "transclusion," "compound document," "open hypertext publishing," "open hypermedia publishing" and "softcopy," were introduced to the computer world at various times over the last twenty-seven years in connection with this project.

1. It is to be a generalized server for compound documents (which we have more recently called "works"), including data portions in any possible data type.

2. It will register and maintain evolving interconnections of all types, even among objects being changed by different parties.

3. Security is non-hierarchical and based on explicit trust boundaries.

4. It will support links, transclusion and versioning, both by authorized authors and those not having permission control (transclusive derivative versions without name permission).

5. A collection of servers functions as a single server; the specific address of the document need not be specified in the protocol.

6. It is intended to scale up cleanly to astronomical sizes-- millions of documents, millions of servers, millions of links and transclusions. Thus a continual design concern has been the interacting problems of combinatorial explosion and feasible response time. Finding internal and external design tradeoffs to make response times feasible on a huge scale has been extremely difficult. (Performance is now designed to scale up in a loglike fashion, so as to be feasible on the contemplated scale.)

7. It may be assembled into an electronic publishing network, maintaining copyright and implementing royalty.

8. The publishing network will support open hypertext publishing, with attachments and variants by non-authorized contributors.

9. Design is based upon a sophisticated understanding of the dangers of the world, and concern for many kinds of attack and falsification at all levels.

THE NEW DESIGN

The resulting new design, approximately the seventh major design in the history of the project,* represents the forthcoming product of XOC, Inc. We shall refer to it as the "1993 version." Its design could best be called "non-obvious."

---

* There have been approximately six previous major designs-- three designed by the author (the third with the help of John V.E. Ridgway and Calvin Daniels), two

a. Unispace. The Xanadu space is first of all a single unified docuverse. One Xanadu server, unconnected to others, presents a Xanadu docuverse; two Xanadu servers, connected through the back-end protocol,† manifest a larger Xanadu docuverse. Thus the division into servers is not logically intrinsic in the docuverse, as it is in so many digital document storage schemes.

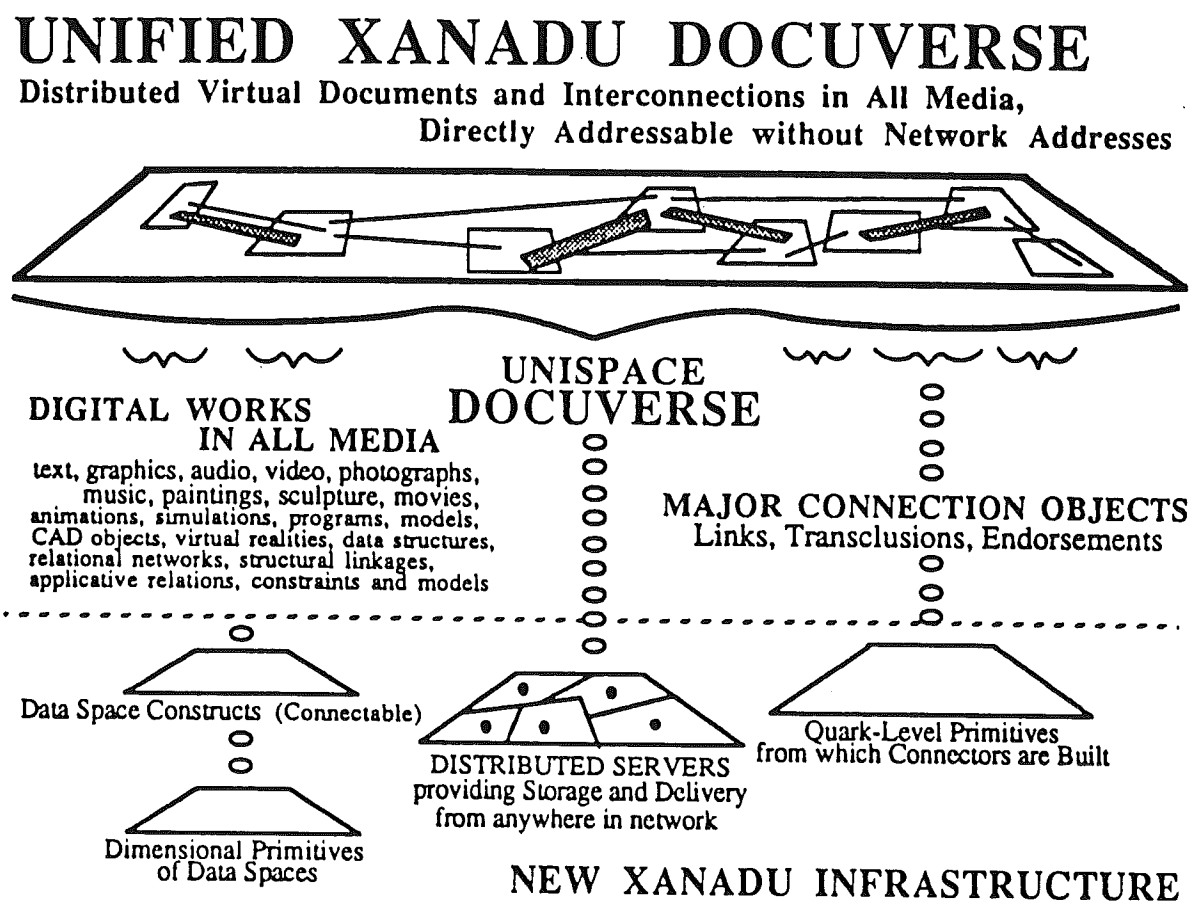

b. The user protocol. Interaction with the system is (as in the 1988 and earlier designs) a dialogue of interactions between the user's screen machine and the server. The server delivers fragments as requested. However, in the new protocol, a loosened system of dialogue has been adopted, such that user requests may

---

by William Barus, and the version described in (1) and completed in 1988, designed by Mark Miller, Stuart Greene, Roger Gregory, Roland King and Eric Hill. The current design is by Mark S. Miller, Dean Tribble and Ravi Pandya.

† Not implemented in the outcoming first product.

continue before the old ones have been answered, and answers are allowed to return without time limit, a necessary feature for the unbounded network we envision.

c. Forms of connection. The principal interconnections of the 1988 Xanadu system, reported in (1), were the generalized link and the transclusion.

The 1993 version adds a new connection type, called an "endorsement." It may best be described as a connection from the name of a user or account to a work or document. The endorsement is untyped, modeless and unidirectional (though like all Xanadu connectors, viewable from both ends). It was originally intended to signify approval, as appropriate to a literary system where recommendations are important.

d. The quark level. In the current Xanadu structure, all user-level objects are high-level constructs. This applies to both connectors and linkable structures. No longer are the fundamental units (such as the document and link) hard-coded into the system; instead, the units themselves are constructs at a technically accessible level. The various constructs of the user-level docuverse-- works, editions, links, etc.-- are built up from lower-level primitives. What we may call the "quark" level is the system by which connectors and other entities are constructed. The system contains several primitives, visible at the user protocol level, that implement these structures.

Foremost of these primitives is the transclusion mechanism. Various forms of transclusion are now becoming popular, such as "dynamic data exchange" in Microsoft Windows, and file aliasing in Macintosh System Seven. But transclusion in our sense, the most general case of virtual inclusion of selections specifiable to the byte level, is not usually implemented; where transclusion exists, it is sometimes implemented as a type of link.* We do the opposite, representing links as built from transclusions. In the 1993 Xanadu system, paradoxically, transclusion is in a sense the most fundamental relation, and transclusion search *the most fundamental operator*, which is filtered and refined by various qualifiers based on endorsement and other particulars.

---

* For example, one hypertext linkage scheme implemented at Sun Microsystems, reported by Amy Pearl, treats transclusion as a type of link, called the "include link".

An unexpected development has been the generalization of the endorsement into this system of construction. Originally conceived as a primitive for authentication and approval at the literary level of the docuverse, it became embraced into the underlevel construction of other mechanisms. Effectively this use of the endorsement reflects the fundamental philosophy of Xanadu design, the fundamental presumption of literature: that *all information has a point of view.* But that is only a metaphysical summary of the practical form of construction selected for our lower-level entities.

e. Data spaces. Text and graphics of different kinds are no longer implemented directly as primitives in the new Xanadu system; rather, they are constructed from the space up. Lower-level Xanadu programmers may define spaces and cross spaces in whatever dimensionality, mapping the structures of interest to them into these spaces. For instance, three-dimensional CAD and virtual-reality spaces are built in this fashion. Because they are based on the Xanadu lower-level construction system, they will link and transclude as well as text and photographs.

f. Security and write permission. Many security systems are hierarchical, with the administrators of the system viewed as incorruptible. Our design philosophy assumes that any administrators may be corrupted and any server nodes or copies of works may be corrupted, and that security boundaries overlap and interpenetrate without hierarchy. The user may view and work on all works or documents in the union of his permissions, or "clubs."

All users have write permission on whatever they see, since any user modification to a locked document effectively generates a new document by transclusion.

BIBLIOGRAPHY

1. Theodor Holm Nelson, *Literary Machines.* $40 (1992 international price) from Mindful Press, 3020 Bridgway #295, Sausalito CA 94965. Japanese translation forthcoming from High Technology Communications, Hachiko Building, 1-34-6 Takadanobaba, Shinjuku-Ku, Tokyo 169; fax 81-33-23-27-396.

2. Xanadu Operating Company, preliminary documentation for Xanadu Hypermedia Server 1993.1. (Not yet published.) XOC, Inc., 1891 Landings Drive, Mountainview CA 94043; Tel. 415/ 903-1300, Fax 415/ 903-1399.