

The Open Society and Its Media

Mark S. Miller

with E. Dean Tribble, Ravi Pandya, and Marc Stiegler

Electronic media present tremendous opportunities for improving the nature of society. I will address how discourse affects society, and how changes in media may improve societal discourse. Then I will describe the Xanadu¹ system, and how it was built to achieve these goals.

Improving Society

Improving society is a difficult task. More generally, improving complex systems is a difficult task. If you cannot figure out which way is up, see if you can figure out which way is down. Doug Engelbart, back in the early 1960s, wanted to explain to people why interactive systems would make a significant difference to their lives, and to their ability to express ideas. The origin on the axis is what people were doing at the time – writing with pencil and paper. When he found himself unable to communicate to people how much better things could be, he contrasted their current experiences with how much *worse* things could be. He tied a pencil to a brick, handed it to people and said, “Okay, now write.” People found it very difficult. The unwieldy nature of the tool interfered with their ability to express ideas. With the pencil and brick for contrast, he effectively asked two questions: “What made the difference?” and “How can we move further in the other direction?” (Engelbart 1962). This experiment showed people how important their tools and their media were to their effectiveness, and helped them start to see the next brick to remove.

Karl Marx performed a similar experiment on society over the course of most of this [the twentieth] century. Karl Marx tied a very large brick to a very large pencil and the last few years have revealed the result to be far worse than the even his harshest critics imagined (Popper

Originally published in *Extropy: The Journal of Transhumanist Thought* 12 (1994). Copyright © Max More.

The Transhumanist Reader: Classical and Contemporary Essays on the Science, Technology, and Philosophy of the Human Future, First Edition. Edited by Max More and Natasha Vita-More.

© 2013 John Wiley & Sons, Inc. Published 2013 by John Wiley & Sons, Inc.

1950). What made the difference between the societies? Two important elements were open markets and open media. How can we move farther in the other direction? In this presentation, I will be addressing the nature of open media, how they differ from closed media, and how social hypertext systems can enhance the advantages of those media. Applying information technologies to the further opening of markets is left as a mission for the reader.

Media Matter

Media matter, because it is in media that the knowledge of society evolves. The health of the process by which that knowledge evolves is critical to the way society changes. Karl Popper, the epistemologist, had the insight that knowledge evolves by a process of variation, replication, and selection, much as biology does. “Variation of knowledge” is what we call “conjecture” – hypothesis formation, tossing new ideas out there. “Replication of knowledge” is the spread of ideas through publication and conversation. “Selection of knowledge” is the discrediting of conjectures through the process of criticism.² The ability of our knowledge to progress over time depends on an ongoing process of criticism, and criticism of criticism. The ideas that survive the critical process tend, in general, to be better than those that do not.

In closed societies, when arguments cannot be spoken, hard truths cannot be figured out. When people cannot openly criticize, cannot openly defend against criticism, or cannot openly propose ideas that conflict with the official truths, then they are left with mistrust and cynicism as their only defense. This leads to the simple heuristic of assuming the official truth is always wrong. For example, because *science* was promoted by the Soviet propaganda machine, pseudo-science is on the rise in Russia. Because anti-Nazism was promoted by the East German propaganda machine, Neo-Nazism is on the rise in East Germany. The official truth is neither always right nor always wrong. Society needs a more sophisticated process for judging claims.

Our society does have open media. Are we in the best of all possible worlds? Are our media good enough? Can they be made significantly better? Among our media, TV is so bad that it is a joke. Only slogan-sized ideas can be expressed. We prize the quality of discourse in our books and journals, but critical discussions in them are only loosely connected. Starting from the expression of an idea, it is hard to find articles that criticize that idea. When arguments cannot be found and navigated, the next harder truths *still* cannot be figured out.

Xanadu

I rejoined Xanadu in 1988 largely because of fear about the dangers of nanotechnology, coupled with incredible excitement about the promises of nanotechnology. In looking at the dangers, I saw that none of us individually is clever enough to figure out how to solve those problems. The only hope that I saw in 1988 – I no longer believe it is the only hope – is that by creating better media for the process of societal discourse and societal decision-making, we stand a much better chance of surviving the dangers posed by new technologies, so that we may live to enjoy their benefits.

I am about to go through the elements of the hypertext system we built. Xanadu has frequently been called Golden Vaporware, and many people have wondered whether this is a never-ending project. One of the things I want to emphasize when I go through all of these features is that I am only referring to the features that are now running in the software. We planned on and anticipate

other features, some of which will be mentioned in the *future plans* discussion, but the body of this presentation will only cover what is implemented and running.

First, I will discuss the four fundamental features – links, transclusion, versioning, and detectors. Marc Stiegler will then present an example using them. Then I will describe the remaining four features – permissions, reputation-based filtering, multimedia, and external transclusion, followed by some concluding remarks.

Links

Hypertext links are directly inspired by literary practice. Literature has many different kinds of links connecting documents into a vast web. Textual examples of these links include bibliographical references, marginal notes, quotation, footnotes, and Post-it notes.

We propose to build engines of citation, so that people can navigate this vast web of literature at the click of a mouse. Most computer text systems are predicated on a misconception that the meaning of a document is represented purely or primarily by its content. Documents are not islands. Conventional computer text systems put their effort into the appearance of individual documents. My experience in reading documents (especially reading a literature with which I am not familiar) is that it is difficult to understand documents without their context. A context helps answer questions such as “What were the ongoing controversies that the author had in mind?” “What views was he supporting or attacking?” “What attacks was he guarding against?” We must understand this whole web of connections in order to understand the documents we are reading. The Xanadu system is built to provide as much support for this contextual information as for content.

With the ability to follow the links in this vast web of documents, is it not easy to get lost? How does one stay oriented? One answer to these questions is *guides*, a new kind of document that provides an orienting view together with links into the existing literature. I expect guides to come largely from people making their own organizing views of a literature and then cleaning them up for publication, so others may benefit from their work.

Hyperlinks

Because “nanotechnology” is now used by many to mean any technology approaching the nanometer scale, we have been forced to retreat to the term “molecular nanotechnology.” Hypertext terminology has gone through a drift similar to nanotech terminology. The Xanadu project is the one that coined the term “hypertext” and originated the notion of the hypertext “link.” However, because the term *link* has come to be viewed as something much less capable than what we meant by it, we are now calling it the *hyperlink*. The distinction between the link and the hyperlink is crucial for supporting active criticism in open media.

Hyperlinks are fine-grained, bidirectional, and extrinsic. Frequently, an argument is not with a document or chapter as a whole. It is with a particular point that someone made at a particular place in the text. For example, someone refers to the fourth law of thermodynamics, and someone else writes a criticism saying there is no fourth law of thermodynamics, linking it to the original. The fine-grained property allows the link to designate the particular piece of text with which one is taking issue. Bidirectionality enables readers of the original document to find the

criticism, enabling them to exercise fine-grained skepticism, and to constantly ask themselves, “What is the best argument against the thing I am reading *right now*?” and then, “What is the best argument against that, in turn?” Links provided by other hypertext systems generally have been only in the forward direction, enabling a reader to find those documents *referenced* by a given document. However, to find criticism, the reader must find the documents that *refer* to the document they are reading.

Extrinsic linking is the ability to link into a document without editing it. Several other systems support the creation of links that are fine-grained at the targeted end, but these others do so only by modifying both source and *target* documents.³ Critics normally will not have the ability to modify the documents they are criticizing. They could spin off their own version into which they attach these links, but then other readers *still* cannot find these criticisms from the original documents.

Part of what we mean by “open media” is that everyone who is connected to the system can read what they are permitted to read, can write new things, and can make them accessible for others to read. This includes making links to anything that they have read, so that anyone else who reads the original can find the material that has been linked to it. All readers of the system are potential authors. We can think of this process as *active reading*. Frequently, people make marginal notes to themselves. This is a medium in which readers can share such things with each other. When much writing is commentary about other text, the commented-on text is the best rendezvous point for the authors and readers of commentary to find each other.

Emergent Properties

This kind of accessible criticism can provide for decentralized consumer reports. When people post on the system documents that are either products or descriptions of products, customers of those products can post criticisms of them. What did they think of using them? This commentary can guide the purchasing decisions of others.⁴

A particular capability we are used to in conversation (one that is almost impossible to successfully attain using paper-based literature) is hearing the absence of a good response to an argument. A reader not only can see what the most compelling arguments are against some statement, but also see when there are none, or when all the seemingly compelling arguments have been successfully refuted. Such absences are quite obvious in conversation. Electronic media can make these absences obvious as well, but in a context where the absence will be much more telling, because the missing argument could have come from a much larger audience over a more extended period of time.

Other hypertext systems with their unidirectional links reproduce the asymmetry present in our paper-based media – it is much easier to find something that a document cites than it is to find those documents that cite a given document. One of the effects of this asymmetry in paper media is the pathological division of scholarly fields into disjoint “schools.” Instead of healthy intellectual engagement, debate, and cross-fertilization of ideas, we see a process of increasing inability to communicate between schools, and more preaching to the converted within a school. The terrible irony of attempting scholarship with unidirectional links is that *the very attempt to engage in healthy debate across schools accelerates the pathological division process*. How does this occur?

Let us consider two schools within a discipline. Generally, students within a school see the documents supporting the positions of that school. The students also see criticisms of

documents in the other school. Intellectually eager and honest students, seeking to know both sides, occasionally will follow these criticism links forward. The result is that they will see the parts of the other school's literature that is *most soundly criticized* by their own school, immunizing them more and more against the foreign ideas. With bidirectional links, these students can also find the greatest challenges to their own school. Bidirectional links also enable them to find the *most telling criticisms* of the ideas they are inclined to accept.

Transclusion

Before there were modern economies, there were many little villages, each with their own little manufacturers having to go through a large amount of the production process themselves. These economies were, therefore, much less productive. An individual baker or shoemaker, for example, would reproduce the same kind of work that was being reproduced in many other villages, and would have to fashion a shoe, not quite from raw materials, but without intermediate goods. In extended economies, people can build on one another's work, and there can be a finer-grained division of labor and knowledge, with better specialization.

Now, with respect to literature, authors are frequently faced with the task of re-explaining and restating background material that has been explained well elsewhere. If you could just borrow that material, those existing good explanations, and incorporate them (with automatic credit where due), your efforts could be spent stating what is new. We introduce the concept of *transclusion* to separate the arrangement of a document from its content. There is an underlying shared pool of contents, and all documents are just arrangements of pieces from that pool. [For example, imagining] three circled appearances of the same text are actually just one piece of text in the underlying shared pool of contents, and it just happens to appear in three different arrangements which constitute three different documents. We refer to the three documents as *transcluding* that piece of text. The separation of content and arrangement also leads to good support for incremental editing. Different versions of a document are just different arrangements of mostly shared content.

This is more than just a hack to avoid the storage cost of making separate copies. Hyperlinks are linked to the content, not to a span in an arrangement. Therefore, when someone writes a criticism of content as it appears in one arrangement, that criticism is visible for the same content as it appears in all other arrangements, including arrangements that were made before the criticism was attached. The normal incremental editing process of a single document is analogous to *evolution by point mutation*. The ability to transclude text from other documents allows the analog of sexual recombination. Were links visible only from the arrangement into which they were made, both variation processes would destroy selection pressures by leaving criticisms behind.

Remembering the Past: Historical Trails

As you are editing, an *historical trail* gets left behind – breadcrumbs in history space. The historical trail is simply a sequential arrangement of the successive arrangements of contents. This is yet another kind of context important for understanding. “Things are the way they are because they got that way” (Weinberg 1985).

Understanding *how* they got that way often aids our understanding of what *they are*.

Preparing for the Future: Detectors

In addition to looking into the past, one also reads a literature knowing it will be changing. How can one keep up? To keep track of what is happening, to keep up with changes, we introduce *detectors*. One can post a *revision detector* to find out when things are edited, when new versions of something appear, and then one can use *version compare* to find out how they are different. With version compare, one can engage in *differential reading* – reading just the differences between the current version and the version most recently read.

Link detectors are a way of finding out when new links are made to existing material. Let us say that you published something, and you want to find out when others post comments on it. You would like to be informed of comments, but you do not want to have to go back and constantly recheck all the things that you have written, so you post a link detector on all the things that you have written, as well as on other documents on which you are interested in seeing further comments. You want to see what people will say about them. As new comments are posted on those documents, you are continually informed.

Email is just the special case where you establish a canonical point in the literature, for each person – a place others link to in order to send that person a message. That person simply has a link detector there saying, “Show me all new things that are attached to *here*.” This generalizes to treating any shared point of interest in the literature as, in some sense, a “mailbox,” or a “meeting room” for further conversation or conferencing about a topic. Canonical documents become meeting places. Should two disjoint discussions about the same topic spontaneously form in two places, anyone who notices can just make a link between them. The link detectors of each community will then inform them of the existence of the other.

At this point, I will shift over [to the work of] Mark Stiegler and Dean Tribble, who demonstrate, using the Xanadu software, an example involving exactly the elements discussed so far.

The WidgetPerfect Saga

This is a tongue-in-cheek story about how a hypertext system was able to save several thousand jobs. One special characteristic about this story is that it is a story from the year 1997.⁵ It is a story about one of the events that took place at the company – most of you have heard of it – called WidgetPerfect. WidgetPerfect is the second largest manufacturer of widgets in the world, second only to their big competitor, Microwidget. The people at WidgetPerfect in the year 1997 had identified a really significant opportunity in the upcoming expanding environment of widget components technology.

They were developing the world’s first fully modular widget. They had a team working on it. Dan Tribble was in charge of the preparation of the marketing materials for the modular widget. Ruth was in charge of the technical work team, and John was in charge of the budget and finance, as well as all the costing. At this point, the modular widget was in prototype stage when a very unfortunate thing happened. Microwidget, the big competitor, came out with a partially modular widget, hitting the marketplace first with an inferior product. It was technically inferior, but nonetheless it was in the marketplace first.

Dan was examining this Microwidget, partially modular widget, and it was overall inferior. Nonetheless, it had one really striking improved feature. It had a funculator made out

of titanium, whereas the fully modular widget that was being developed by Ruth only had a duralum funculator. This was an important improvement for certain key market sectors. Even though the partially modular widget did not have anything comparable to a thermoplastic coupler or a hypervelocity rotator, they had to make a change.

So, Dan created a new document in the marketing requirements describing this titanium funculator. He attached a link to the part of the technical plan that specifically referred to the duralum funculator in the current plan. He made that a new requirement.

Now, Dan knew that in order to get anything to happen with improving the widget prototype, he would have to talk to Ruth. He was reaching for the telephone to call Ruth when Boeing, the largest purchaser of widgets in the world, called him about a \$15 million widget order. He got distracted with this purchase, and he never quite got around to calling Ruth.

We have good news. Ruth, knowing that the success of her technical design depended on her being able to respond promptly to new requirements, had attached a link detector to her technical plan. This link detector would be constantly watching for new links of the link-type *requirement* to be attached. When Dan attached the new requirement to the duralum funculator, Ruth's link detector went off. Ruth was alerted. She followed the link detector out to the link, followed the link back to the new requirement, saw what the required change was, and modified the technical plan to reflect the use of a titanium funculator.

Well, this was all very fine, except for an additional problem. As I think everyone here knows, titanium is considerably more expensive than duralum, and so this had some significant effect on the manufacturing cost. Ruth knew that this would have an impact on the budget, and she was reaching for the telephone to call John when smoke started billowing from the laboratory where the prototype of the modular widget was being manufactured. She ran off to deal with the emergency and never quite got around to calling John.

We have good news. John, knowing the success of his budget was completely dependent on his responding to modifications to the technical plan, had attached a revision detector to the technical plan and this detector was constantly watching for updates. So, when the technical plan was indeed updated, John's revision detector went off. He followed the revision detector up to the technical plan, used the hypertextual version compare capabilities based on the transclusion relations, compared the new version of the plan to the old, and found that the change was deleting duralum and replacing it with titanium. He then went back into the budget and updated the budget documents to reflect the increased costs caused by the use of titanium.

As a consequence of this, the modular widget program was completed on time with a fully adequate specification. It was a completely superior product. It blew Microwidget off the face of the Earth. As a consequence, thousands of jobs at WidgetPerfect were saved.

Permissions

A social system, to a large extent, is a system of rights and responsibilities. Xanadu has an extensive permission system called the *club* system, intended to deal with some of these issues. Bob has sent it as a mail message to various people in a blind carbon copy ("bcc") relationship. Alice and Chuck are both members of the *bcc* club of people who have permission to read this document. Bob, though, is the only member who can read or edit the *bcc* club. If this were a *cc* list, Bob would still be the only person who could edit it, but it would be self-reading. Everybody who was a member of such a *cc* club could see who else was a member of that same club.

This demonstrates a principled answer to permissions *meta-issues* – one can distinguish between who can read a document, who can read the list of people who can read a document, and who can read that list, out to any desired degree of distinction (and similarly for the editing dimension). However, infinite regress and needless complexity are avoided by using clubs that are self-reading or self-editing (or both) whenever further distinction is currently not necessary. Should such distinction later become necessary, it can always be introduced by someone with appropriate edit permission to the club in question. Users only grow meta-levels on an as-needed basis.

Our permission system also supports the notion of *accountability*. All actions in the system are taken *by someone*. When you look at information in the system, you see some identity attached to the actions taken on the information. There are no official truths. There is only who said what, and the structure of the system reflects that.

Reputation-Based Filtering

One of the potential pitfalls of an open hypertext system is the junk problem. The ability to find good commentary and criticism will be especially important when reading very important documents, but it is precisely on *these* documents that one expects to be inundated with tons of worthless or irrelevant links. Without a filtering mechanism, it would be on exactly the documents for which one most needs good commentary that the provision of commentary would be most useless. For example, imagine how many links there would be onto the First Amendment to the Constitution.

Links can be *endorsed* as worth reading by various readers. However, no one may endorse with the identity of another. Different endorsers will establish varying reputations with different readers, much as with movie reviewers. Readers can filter their views of links into a document both by *who* endorsed as well as by link-type. When even this mechanism gives too coarse an answer, one can rely on documents such as a hypothetical *Guide to the Citations to the Bill of Rights*, endorsed by a reputable publishing house. This very same link-filtering ability is also what allows one to find such guides in the presence of a swamp of links.

Hypertext + Multimedia = Hypermedia

Increasingly, ideas are being expressed in media other than text, and increasingly, computers are used to handle these other media. We usually refer to *hypertext* because text is the most important case and the clearest example. However, nothing I have presented, none of the things you have seen the system do, is in any way specific to text, or even to media that have linear flow to them. It all applies equally well to a variety of other media (such as sound, engineering drawings, Postscript images, and compressed video). In all cases, one can make fine-grained links, edits, transclusions, and version compares (even if the data is block-compressed or block-encrypted). Although the implementation has some optimizations targeted at text, in no way does the *architecture* make any special cases for text. Documents can, of course, be composite arrangements in which several media are mixed together.

External Transclusion

No software system is an island. We do not imagine that once the product is available, everyone will instantly take all information to which they want access and transfer it into Xanadu. We have to coexist with many other systems for many good reasons.

We handle that with *external transclusion*. Our documents are able to transclude into arrangements that are within the system. These, in turn, are able to represent transclusions of materials that are stored elsewhere. By perceiving other systems through the window of Xanadu, you can see those other systems as if all those documents were within the Xanadu system. Through Xanadu, I could follow a link from a WAIS document into a Lexis document, even though neither system has any notion that such a link even exists. It is not just that the Xanadu system is not an island that we have to coexist with everything else; it is that through Xanadu, *those* systems are able to coexist with each other in a way they are unable to now, making *them* into non-islands.

Conclusions

When we started building the system, we were thinking purely in terms of paper-based literature – of writing. What we have built is something that has many of the best aspects of both writing and conversation. Many of the aspects of each are complementary. Many virtues of conversation make up for flaws in writing and vice versa. We found ourselves building a system that supports the dynamic give-and-take of conversation, and the persistence and thoughtfulness of literature.

Our status is that we currently have a working, portable server. It has some bugs in it, including some performance bugs, but we are working on it. However, all the features that I talked about so far, work. We are continuing ahead with the effort on both the server and the front end. The front end is in a preliminary stage. We consider it adequate to show that the server is real, and to exercise its features. We plan to do a much better front end. The protocol between the front end and the server is very stable, and has been stable for a long time now. Our plans are to get investors, and to finish both the front end and the server. The target for our first product is small to medium-sized workgroups within companies that have a large body of documents they need to be managing and evolving.

Our first product lacks one major feature. We provide hypertext because documents are not islands. We make the system interpersonal because people are not islands. We provide for the transparent windowing into other systems because no product is an island. However, for the moment, each server is still an island with respect to the other servers, and so each workgroup is also an island. We have designed the system so that, soon after first product, we will be able to weave all the servers together into a transparent distributed system. When you follow a link from one document to another, if the other document is not here but in some server in Tokyo, it will be transparently fetched for you, and the only thing you will notice is that following that link took longer.

For any media to radically improve the process of opinion formation in society, we believe it *needs* features equivalent to fine-grained, bidirectional, extrinsic, filtered links. These links must not get lost when the documents to which they are attached change. Issues of authority, privacy, and responsibility must be handled in a robust and secure fashion. Open entry of readers and editors is crucial for open discussion. Open entry of server providers is less obvious, but equally important, in order to make centralized control impossible. We will be providing support for people who want to do online services based on our software. All of this is necessary to achieve our open electronic publishing dream. In so doing, we hope to improve the quality of public debate, in order to obtain the benefits of the open society yet again.

Notes

We thank the whole extended Xanadu team for having struggled together for many years on a project that has been at least as much a cause as a business. We thank Eric Drexler for exploring the relationship of hypertext publishing to evolutionary epistemology (Drexler 1991). We thank Anita Shreve for extensive help in editing this presentation.

- 1 The Xanadu trademark has since become the sole property of Ted Nelson.
- 2 Karl Popper originally proposed that selection proceeds by a process of refutation. See Popper 1959. His student, William Bartley, generalized this to criticism. See Bartley 1962.
- 3 Examples include World Wide Web *anchors*, Microsoft Word *bookmarks*, Lotus Notes, and Folio Views *Popup text*.
- 4 The use of bidirectional links for decentralized consumer reports is already happening on the American Information Exchange.
- 5 This essay was written well before 1997, thus the fictitious *tongue-in-cheek* story is actually a hypothetical scenario about electronic media.

References

- Bartley, William W. III (1962) *The Retreat to Commitment*. Chicago: Open Court Publishing.
- Drexler, K. Eric (1991) "Hypertext Publishing and the Evolution of Knowledge." *Social Intelligence* 1/2.
- Engelbart, Douglas C. (1962) "Augmenting Human Intellect: A Conceptual Framework." *SRI Project 3578* (October).
- Popper, Karl R. (1950) *The Open Society and its Enemies*. Princeton, N.J.: Princeton University Press.
- Popper, Karl R. (1959) *The Logic of Scientific Discovery*. New York: Harper & Row.
- Weinberg, Gerald M. (1985) *The Secrets of Consulting*. New York: Dorset House Publishing.