Appendix A
1958 Schematics Paper

Ted Nelson
Social Ph.
May, 1958

Schematics, Systematics, Normatics.

This is my own system. I think it makes sense, and hope the sense to be communicable. If an apology is called for, I will save it for more careful, lucid and thorough presentations of a later date, unless I am encouraged to desist.

There are three ideas here, one for each section. Since I have pulled them out of the air, I would rather someone else had said them first than that they were useless.

I have not got the slightest idea what and how the herein relates to the main stream of philosophy, nor to symbolic logic or statistics (- although I feel a strong feeling toward factor analysis.) It has been more enjoyable to blunder in the dark than to grope for the light switch.

However, I think I do know how this relates to structual linguistics; I would like to believe it is the sweited inflation of the "linguistic model" to universal proportions. Such titanic ambitions aside, I would be most pleased if anyone finds it useful.

PART ONE. SCHEMATICS

There are two principles in description which I find titillating. One is that the same thing will often be described quite differently by two people. To illustrate pictorially. You could draw any of the following curt sketches, and call it a "man."

(1)

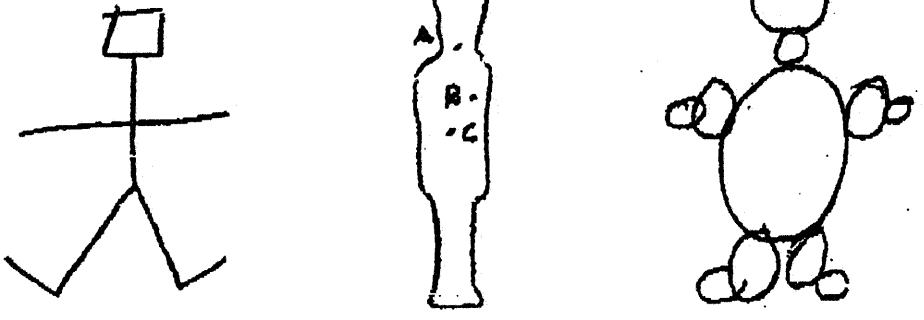

The other principle is that the same description will often fit a number of different things. Again, the pictures.

(2a)

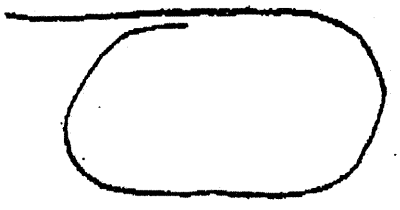

Path of marble in toy;
path of evasive rabbit;
path of foot race;
path of particle in cyclotron.

(2b)

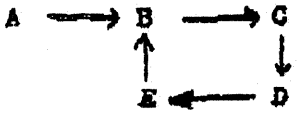

Feedback; sociometric layout; Arthur Murray step.

It seems to me that the two principles merit attention under a single heading. Let me introduce the term "schematic" for <u>any</u> description or description, not necessarily a picture. I do not use the word "description"-- or, for that matter, "concept"-- for several reasons. 1) A description can be "good" or "bad," "complete," "incomplete," "true" or "false." I would not like to say

[illegible] could hold for the schematic, although it will certainly be that the application (or correspondence) of a/schematic given to a given context will be good, bad, true, etc. 2) A description or a concept, to exist, must be made, held, stated or thought. A schematic need not. A schematic, like a number, might be said to subsist, like a number, no one need ever have thought about it. 3) I will later discuss the transformations among schematics; these would be unthinkable for "descriptions."

A schematic, unlike a description, need not be explicit; while we can reverse it and say that a description is an explicit schematic. It is important, though, that a schematic be explicable; that is, that it may be set forth in communicable, examinable and delimited form. But no matter how it is set forth, it is still the same thing: a set of things and relationships arbitrarily abstracted and differentiated. Arbitrarily, of course, does not imply limitless option. Rather, the sensible rules for reasonably forming schematics in various cases and places are diverse.

The schematic is frequently encountered as a figment of education, where a crude collation of shapes and vectors is keyed to such things as "Roman Empire," "God," "realism." These can be, however, sensible ways to think and talk, and so I would like to examine aspects of this way of talking and thinking.

Often a schematic may be characterized by someone as being "fundamentally this-and-so," or "generally this-and-that." Is it possible or practicable to characterize schematics in this fashion? Here is a list of terms which have been applied to schematics: [illegible] or Manichean, [illegible] or, for that matter, "concept" [illegible] Dyadic, or dichotomized, or pluralistic, "and," "separate," "complex," or "many-pluralistic." I [illegible] to say

d. Circular.

e. Dynamic.

f. Holistic.

g. Coherent. Or unified.

e g. A monistic schematic would be one possessing an overriding single principle:

(3) Everything in the universe is ordained by God.

The sentence immediately calls forth the suspicion that the single principle does not say anything and can be cancelled out. This would be fun, but I fear it is not so. Let us take (3) above, and dissect it.

(3a) The physical laws of the universe was ordained by God.

(3b) The particular circumstances pertaining in the universe were ordained by God.

(3c) The acts of man are ordained by God.

We could reduce this to:

(3a') The physical laws of the universe, the particular circumstances pertaining in the universe, and the acts of man--

(3b') -- are or were ordained by God.

We wondered whether we could cancel out the "God" part, or whether it would be illicit. Suppose we try it, correcting for grammar and suggestion.

(4) There exist in the universe physical laws, particular circumstances, and acts of man.

This is a perfectly reasonable position. However, by inspection we know that it is certainly not the same belief as (3), (3a,b,c) or (3a',b'). One obvious reason is that it yields none of the propositions (3a,b,c). We note that (4) resembles (3); but then, so do

(5a) The physical laws of the universe were ordained by God.

(5b) The particular circumstances were ordained by God.

(5c) The acts of men are acts of Free Will.

b. The dyadic schematic would rest on two principles or things (say, "God and the Devil are jointly accountable for everything of importance," "All behavior contains two aspects: the utilitarian and the symbolic." Such schematics contain the fundamental dichotomy

(6) A & B (alternatively, A versus B, A among the B, watch out for the A in the B, which side are you on? etc.)

Other schematics which might be related to dichotomies would not contain the dichotomy, as: "Don't bother with the in-between ones, just throw them away," "The Negroes are getting lighter and lighter," "There used to be only two kinds of women-- now there are three."

c. The complex, or quite pluralistic, schematic, is one that just doesn't factor easily. E.g.,

(7) The sparkplug fires, and the gas explodes, and the explosion drives the piston, and the crankshaft turns, and the doohickey waggles, ... and it starts over again.

Can we measure, then, how complex a schematic is? Certainly not by a numerical count of the entities or principles in it.

Suppose I were to say, "History consists of the actions of Adam, Eve, Cain... me." If all that mattered were the number of entities, here would be an enormously complex theory; yet it can plausibly be considered monistically, i.e., "History is the actions of all men, ever." So I think we must conclude that there is no measure of how complex a schematic is, except in accordance with how well it fits a paradigm (q.v.) In other words, if I say, "That is a complex theory," (and am not merely indicating thereby that I do not understand it,) I mean a) "there is not some 'obvious' way of factoring the theory-- as 'monistic,' etc."2 -- and b) "the other theories I have in mind to cover this ground do not relate easily to this one."

d. The circular schematic is not merely one whose definitions are circular, since in pure schematics all definitions are circular (see p. Rather, the circular schematic would be exemplified by the children's game "Stone beats scissors, scissors cut paper, paper wraps stone," the tigers chasing around Little Black Sambo's tree, a merry-go-round, etc. It is clear that there is an underlying Circular factor. Obviously-- and this is why I brought it up-- if you can have circular schematics, you can have square schematics, pentagonal schematics, cylindrical schematics, laryngiomorphic schematics (as in acoustic phonetics,) homuncular schematics. Whether they are useful and true in a given context, depends.

e. The holistic schematic does not exist. Actually it is not the schematic which is holistic, but its advocates. A man will claim his schematic to be holistic in opposition to another's.

What the holist contends is that his schematic reckons with some important fact about the thing schematized, and that the other fellow's does not. He may be right or wrong.

f. A dynamic schematic is one defined in terms of its internal relationships, and for which operational definitions are a) not submitted, or b) inordinately impractical. When an operational definition is not submitted, it may be oversight; it may be the as-yet-inadequate formulation of a valuable insight; it may be chicanery. Operational definitions would be impractical when the relationships were unique for every individual and it were impossible to study thoroughly a samplying of individuals. If, in the Freudian scheme, the relation of the Ego and Id is strictly unique, and we must help the patient rather than toy with him, there is no time for definition and experimentation. There are actually several if's here.

Freudian ifs and buts aside, it may often be altogether justified to postulate a dynamic entity. If we see something unified but the light is dim, if we sense an organized change, we have a perfect right (bounded, as rights are, by discretion) to talk about it as a "thing," not pending the establishment of causal and interlocking relationships, but expecting it.

g. A coherent (or unified) schematic has no contradictions. This is, in fact, no more than to say that it is worked out. Let me illustrate:

(8a) I am the Master of my Fate.

(8b) I am the Captain of my Soul.

(8c) The trouble is, people keep pushing me around.

is an obviously incoherent schematic; if the statements were reduced to mutually compatible language, the schematic would immediately engage itself in logical contradiction. However, some quick cogitation will straighten things out. For example,

(8a') If people weren't so rotten, and the situation so unjust, I would have more freedom, and truly be Master of my Fate, Captain of my Soul.

Or

(8a'')
I am the Master of my Fate.

(8b'') I am the Captain of my Soul.

(8c'') I am the King of the Universe.

(8d'') People think they push me around but actually I just went them to think that. What really happens is I order them to by telepathy.

I trust this example will support my belief that "coherence" is actually merely the completion of remarks toward a schematic. I believe, further, that a viewpoint can always-- though with great difficulty-- be made coherent in the face of staggering contradictory evidence; and that if we adopt a particular schematic for a given context, it is merely because proof on our side is simpler. For example, it can be maintained that the earth is flat without denying any evidence to the contrary, through postulating collapsible dimensions and soforth. Philosophers might contend this idea, but in any case I shall not elaborate on it here; all would agree such a strained position to be a ludicrous posture.

After these general remarks, let us seek a clearer way of setting forth a schematic. There are a couple of ways we have used already: the diagram:

(9) (after Reich)

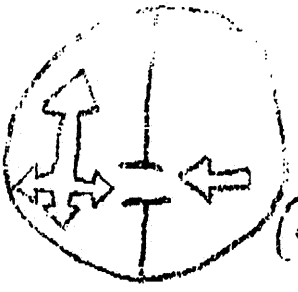

(after Lewin)

This is all very well, but not every schematic may be usefully, conveniently or comprehensibly set out as a diagram on a flat sheet of paper. So we have the set of propositions:

(10a) All men are created equal.

(10b) They are endowed by their creator with certain inalienable rights.

(10c) Among these are 1) life, 2) liberty, 3) the pursuit of happiness.

(One method we have not used here, of probable future value, is the infant technique "Discourse Analysis," under development by the U. of Pa.'s Zellig S. Harris and documented in the journal Language. It is the aim of the method to break down speech or text into its unique isolable components; the end of the operation involved is the formula presentation of a given discourse.)

I would like to try a fourth possibility here, the subsumption of all schematics under the general taxonomic-componential terms A, B, C, D... ("the market has a vegetable counter, a meat counter, a canned goods counter...")

Looking at a schematic "pure," that is, regardless of what its terms mean in common usage, we may regard the parts of a schematic as defining each other. In other words, in the schematic, "John loves Mary," we know, regardless of the English meanings of the words, that "John" is that which loves Mary. Or we might emphasize that "Mary" is being that which John loves. Or, again, it is clear that "love" is what John does to Mary. Ignoring the commitments

of meaning, we can regard the schematic as a set of things the parts of which relate to each other. (Some relationships in the schematic are important and some incidental-- like "John" and "Mary" each having four letters. These we leave out.) If we call our fundamental entities A and B, we can set forth a preliminary statement: "We are talking about John and Mary."

(11a) a, b.

Now we add to the schematic a relationship between John and Mary: "John loves Mary." The notation is my own.

(11b) a*b.

(The relationship could as well be written b*a. The reciprocal, "Mary loves John," and the relationship "Mary and John love each other," must be written in the same way. Means of distinguishing them will be discussed later.)

Now let us add another relationship: "Mary knows that John loves her."

(11c) b**a*b.

Next: "John knows that Mary knows he loves her."

(11d) a***b**a*b.

Each term of the schematic, (11a-d) , represents a <u>part</u>, that is, some relationship worth including in the schematic. The total schematic, as we have set it forth, has five parts:

(11) a, b, a*b, b**a*b, a***b**a*b.

How many possible relationships are there? Obviously an infinite

number, though clearly most of them are of no use whatever.

Now: th specify a schematic is to state all the parts of a schematic; to specify a correspondence between two schematics is to specify all the equivalences betweenm parts. We will try setting forth a schematicm in two ways.

(12)

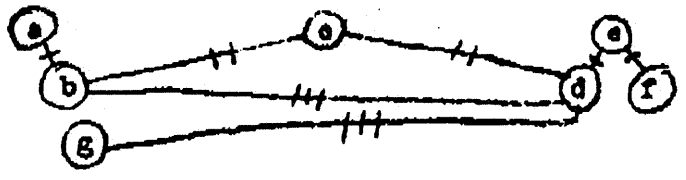

(13) a,b, c, d, e, a*b, b*c, b*d, c*d, d*em, d*g, e*f,
(a*b)**(d*e)-(e*f) (--SEE NOTE,) b*c**c*d, b*d**d*g.

(NOTE: when a part is the relation of more than two relations, it is economical tc use hyphens and parentheses.)

Anything will fit thir schematic whieh has discrete parts corresponding to the parts of this schematic. For instance,

(14)

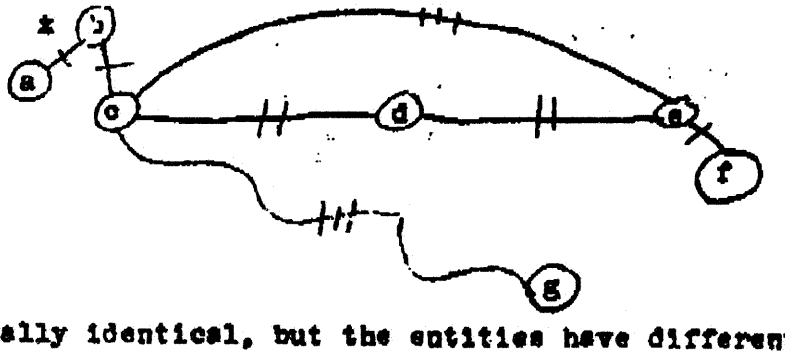

is topologically identical, but the entities have different names.

(One immediate problem that manifests itself is the question of the same schematic set forth so that relations were now entities and vice versa. Obviously, rules of transformation and rotation are called for. I don't know them.)

In any case, the order of a schematic is the number of removes at which relations continue to interlock, that is, the number of

times the relations-to-relations continue to be important. This is probably a function of the number of asterisks in a schematic. I presume that, regarding arithmetic as a schematic, its schematic order would be infinity; so that, as a science evolves, we want to maximize the number and order of true statements we can make in it-- regarding the entire science as the schematic, and each statement as part of it.

The reader may recall the first drawings which all represented "a man," and the curlicue which could be so many things. In both cases, an interesting feature was the resemblance between a number of schematics. In the first instance, both were schematics of the same thing, but we knew this from an outside source (the text.) In the second case, all were similar by inspection.

Let us find a parallel in academic life. In the first case, two different investigators find different ways of describing the same thing-- we find this in psychology: S-R versus Gestaltists, Freudians versus factor analysts. In the second case, someone finds a general principle to hold in several scattered fields: the study of "populations," in demography, ecology, economics, astronomy and microphysics. This shows, I think, that we will want to know when two principles are alike.

Well, when may we say two schematics are the "same" schematic? It is an inadequate criterion of the similarity of two schematics that they seem alike, since human perception, though incisive in some areas, is in others renowned for its lack of subtlety.

Certainly when the same mathematical formula applies in two different fields, we would not venture that it is a different formula because of it. This gives us a first criterion for sameness:

it we will immediately regard two things as "the same" if they are set forth identically in some known arithmetical system, geometrical system, or logical system.

However, outside of mathematics, it is unlikely that the same principle will be formulated in two different fields in such manners that they will seem directly comparable. I would like to suggest that the "sameness" we are looking for is actually optional, or nominalistic; and I will try to show it. I think I can explain more clearly with illustrations than with the new notation. Let us take a schematic.

(15) 

Then let us take several other schematics, and see if we can find correspondences:

(16) 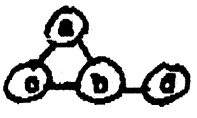

(17) 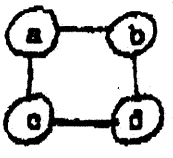

(18) 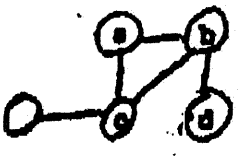

(19) 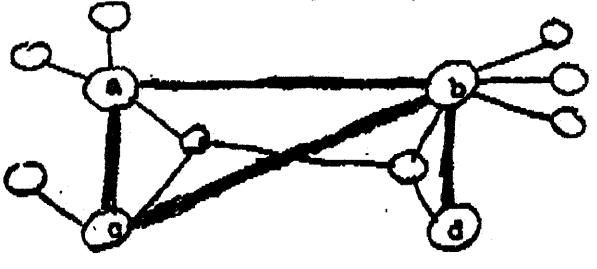

(20) 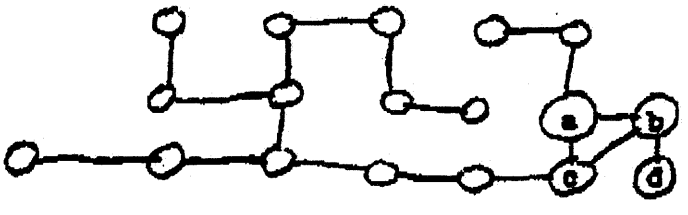

It is clear that all the diagrams have some resemblance to (15). It is hard, in each case, to say what kind of resemblance, at least as far as the real world is concerned; we don't know what the things are, and can only speak schematically. Obviously ~~it~~ (16) is really the same thing, since we can bet that the right angles were incidental. We can't be sure, though; perhaps they meant something. It depends on the schematic's meaning, which we don't know. Or care about, right now. (17) is different, but we don't know how different. We don't know, and cannot ask, whether the diagonal connection between c and b is crucial or trivial. In fact it is neither, since this is just a schematic. All we can say is that one connection is different. Similarly, in (18), we can only say that one thing has been added. In (19), we can see the original schematic (15) as a framework for the whole pattern. In (20), it is hiding off on the side.

This discussion has shown us, I think, that the resemblances among schematics-- hence, among statements, concepts, descriptive systems-- is in some important senses an optional matter: what schematics we see to be alike are alike. They are (of course) sometimes more alike, sometimes less alike. A relation in one may be absent in another. The schematic we seek in another may be a fundamental framework, an accessory pattern, an accidental. Yet it is there if we specify a part-to-part correspondence between it and another schematic.

What about the real world? I would like to say that we can "see" anything in the world in terms of any schematic. What this involves is stipulating a correspondence between a schematic and the world (or something in it,) then specifying the correspondence, part-for-part, as far as it seems to hold true. Thus, similarity.

Whether two things are the same depends on the range we will allow! Except in metaphysics, however, we do this in conformance with empirical indices. Sometimes the alikeness is only a small-scale insight, sometimes a fullblown answer to a theoretical impasse. Regardless of what objects we see, we will often see resemblances-- trivial, interesting, important, overriding-- which move us to say, "these two are similar." The resemblance may be between a barrel and a bottle, between a bubble and a box, between birdwatching and Budhism. In any of these cases, it is impossible to say whether there is "really" a resemblance. Instead, we are comparing two schemàics of the things discused, and finding a minor, major or indifferent set of parts in common. These schematic parts may be intrinsic or extrinsic, internal or exterior: barrels and bottles are both round containers, bubbles and boxes both at least contain space, birdwatching and Budhism at least have names in English which begin with the same letter. This is grossly farfetched; yet it is genuine grounds forassociation and comparison. Whether the comparison is scientific or not depends on a multitude of further things. This would then be a nominalistic theory of isomorphism: that two things are alike-- classif.able in the same way-- on an arbitrary basis.

However, arbitrariness is not wantonness. There are principles which govern our perception, comparison and naming. If two things are alike on a lot of different counts-- that is, they all generally conform to the same schemàic-- then they are the same thing. If, however, we want to discriminate between two things and make two schemàics for two different classes, then they are two different things.

Let me here introduce the word paradigm, as meaning "a

schematic to which another schematic is compared." When we are comparing two schematics, it makes no difference which one we call the paradigm. When we compare three or more, it will be quite important which we call the paradigm.

When we compare other schematics to a paradigm, let us call these its paradigm series. For paradigm (15), we have seen variations in its series where a number of different things happen: where there is slight distortion (16), some of the relationships are not present (17), where something has been added (18, 19, 20). Suppose we set out (15) in taxonomic notation:

(21) a, b, c, d, a*b, a*c, b*c, b*d.

Now we can let a letter stand for each part of the paradigm ("consecutive notation"):

(21') A, B, C, D, E, F, G, H. Where E = a*b
F = a*c
G = b*c
H = b*d

We can call each of these terms, or parts of the paradigm, a feature of the paradigm, since in the whole series we have stipulated correspondences for each one, except for (17). In each case we can say the schematic fills the consecutive features with some leftover (entities and relations not a part of the paradigm,) except for (17), which lacks G. This could perhaps be best called a "null example" of the feature G, since G exists through the rest of the paradigm's series.

If we wanted to quantify the variations of the series from the paradigm, it is quite possible that information thery would give us good indices of "how far apart" two schematics are, by showing how much information would be required to state one in

terms of another.

There are times, alluded to previously, when we want to say that a relation between two entities is not the <u>same</u>-- is less than, or more than, a feature of the paradigm, or is just "different"; this incomplete or irregular correspondence of parts can be in some way annotated so we can tell that there is a difference:

(22) "John loves Mary."

a, b, a*b
Consecutive paradigm:
A, B, C

"Mary loves John."

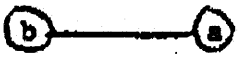

a, b, a*b
How it fits:
A, B, <u>C</u>

As I have said and made painfully clear, the rules for transforming schematics one-to-the-other, and stating them in terms of each other, are beyond the scope of this paper. At best, I am proposing this system as a way of thinking of thinking about things. Once the transformation rules are known, we will have a sort of meta-metamorphosis.

There are a few other points. One is the question of how to select a paradigm for a series if you haven't got one-- which brief will best summarize them all. This is probably a matter of game theory: which one will make the least trouble in the long run, with the fewest null features and qualified fits. Once concept that would probably be useful is the <u>allodigm</u>: a subtype of the paradigm, wherein alternative sets of relationships existed. For example, if the series contained paradigms of two general types: "a, b, c, d, a*c-d, b*c-d, a**a*c-d," and "a, b, c, d, a*c-d, b*c-d, a**b*c-d," the two allodigms of the paradigm would be "A, B, C, D, E, F, G" and "A, B, C, D, E, F, G'".

Another point is that a schematic can be at the same time part of another schematic, just as the same proposition may be a part of two greatly divergent world-views. We saw that in (19), which contained another schematic besides paradigm (15):

(23)

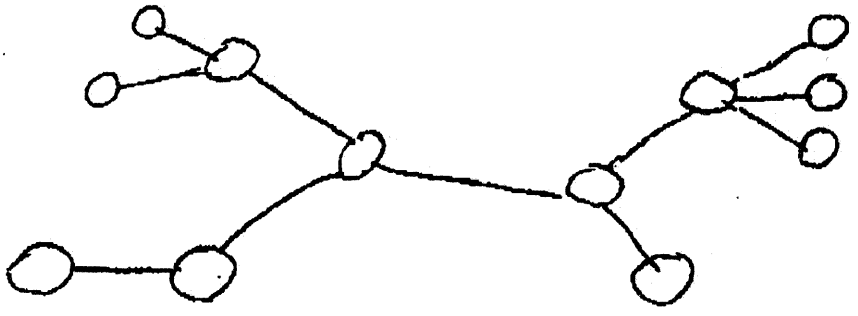

I would now like to turn to a more prosaic topic, that of "meaning." I have already claimed that anything satisfied a schematic if there has been a stipulation of correspondence and a specification of correspondence. In other words, "The Universe is an Elephant" can have a sort of miasmic truth if we schematicize an elephant as having trunk, ears, tail, etc., and explain parabolically how "The trunk of the Universe is scientific law, the ears are the beliefs of man, ..."

We notice that if "elephant" can be thought of as a sort of vague schematic, so can other words; so can principles; so can metaphors; and so on. "Vicious circle," "tight squeeze," "victim," "keystone," "turnbuckle," "melting pot"-- all these carry schematics as part of their meanings; they are principles looking for contexts. However, there are other words with far more anchored meaning: "footpound," "leucocyte," and such. The scientific mind

is impatient with talk whose meaning is not so hardened. Even in science (especially social science) we may espy a principle but be still short of formulating it, quantifying it, operationalizing it, anchoring it. In this case we should not bake it too soon. Outside science-- in home and office and garage-- we have no need for more precision than will get the work done, as, "Hand me the thingamajig," "Watch it!" And in poetry and literature, the anchoredness may be greatly loosened: "Life is like a coffeepot-- you tilt it part way and you get what you want but you tilt it the rest of the way and the top falls out and knocks your cup over" needs no apology, since the schematic correspondence between "life" and "coffeepots" comes clear.

One thing I have not attempted to cover in this paper is the problem of segmentation-- how a "whirlpool" gets divided up into entities and relationships, and how (2a) ever got to be (2b). It would seem sometimes to be like finding a mode or a midpoint, yet often in many more dimensions; how we might chart the differences between two segmentations of the same thing (as in illustration (1), especially the second "man".) It would seem a bit knotty.

PART TWO. SYSTEMATICS.

Having taken a nominalistic stand on the resemblances between schematics, I will now seem to reverse myself and discuss systematics, the study of paradigms and differences of "similar" things in the real world.

A systematic is a real schematic, and far less delimitable. In fine, a systematic is a thing, with its ranges of changes. It will be recalled that my account of a "thing" denied that "things" existed, saying rather that they were perceived, and with some arbitrariness, lumped together in a single class at the will of the namer. This is schematics. A systematic is the occurrences of a schematic in the world, and systematics is then the study of the features that unite the class, and what variations they may undergo while remaining within the class.

Let us take all the cows in the world; that is, all the things that would be referred to by the word "cow," whittling off its fringe connotations. The composite picture they will yield will be the paradigm, "cow." The features of that paradigm will be the things all cows have in common and the lasting, common relationships: those which run through all instances of the class, and which may be sensibly chosen as features of the paradigm: the factors of ability and personality of cows, as they might be uncovered by Thurstonian analysis of reams of tests.

The feature of a systematic, like that of a schematic, need exist in all instances or a definable set of them: that is, hair on top of the head is a feature of the paradigm "men", but in the systematic of "all men" there must be an allodigm for men who are bald.

1. An example: the frisby systematic.

It may be seen, interestingly, that in the game of Frisby an individual has a characteristic thrust, which has personalized aerodynamic qualities; the direction and curve of path and the pattern of flutter vary personally within certain investigable limits. Let us examine the thrust of "B.P."

The flutter, which I have exaggerated, has a characteristic spacing relative to the entire trajectory. Presumably these personal touches are caused by subtle characteristic muscular relationships of the individual. The systematic of B.P.'s frisby thrust, then, is the range of possible thrusts and the consistent formulae that can be advanced to summarize this range. It might be a part of the formula, for example, that B.P.'s thrust fluttered during the second third of the flight (this contradicts the drawings; but no matter) and that the flutter was in direct proportion to the distance; but it would scarcely be that simple.

The brothers "S.", by contrast, throw the frisby in the following ways:

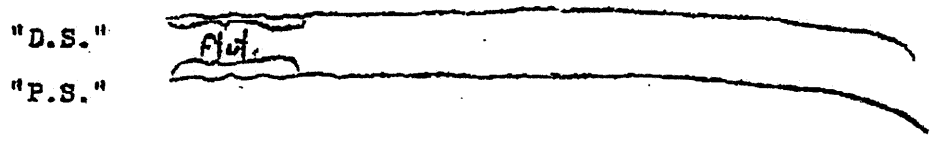

Again, a formula could be developed.*

* Postulations of a "frisby gene" are premature. We wait on studies of identical twins reared apart.

Now: the composite systematic for all frisby-thrusting would be a master formula for all existing personal patterns, showing the correlations and patterns of variation. In other words, it would be a factor analysis.

This is not a dazzling viewpoint; but I would wish to say that some things not now so studied can best be studied in this systematic fashion: as variations-from/instances-of a central pattern.

2) The study of attitudes, values and beliefs, a cavernous cranny of social psychology, has proceeded along arbitrary dimensions in investigation and measurement, dividing "values" (for instance) into "economic," "political," "social," "theoretical," "aesthetic," and "religious;" then devising scales which best measure individual variation along these dimensions. The same has been done for "authoritarianism," a personality-attitude trait created on the drawing board over a protracted and expensive study, "anti-semitism" (a part of the same study,) and many others. I think that these studies are going at it the wrong way. I would submit that a considerably more satisfactory way to study values is by the systematic method, whereby the persistent relationships are uncovered: factorial studies, etc. These would yield some general knowledge of the norms, major contours of belief, major patterns of variation, minor variations, etc. The major patterns would be "points of view," "value-systems" and such, the minor patterns would be significant lines of dissent. I am confident that these patterns would correspond in some important way to major lines of argument, rebukes and compromises. Now, it could be objected that the "values" and other studies were systematic, working from a central paradigm of six dimensions; but I would contend that they are not: no good reason is advanced for

the acceptance of these dimensions, while the "systematic" approach would at least give you major contours of belief as well as a base for making scales of accession to major patterns. Moreover, the six-axis method would only be plumbing inaccurately the same thing we wish to sound systematically.

It will be noted that the majorm patterns of viewpoint would change from population to population, from place to place, from time to time: for this I would like to call them "idionomic," meaning that they cannot be expressed according to an imported schematic. Therefore a given set of opinions will factorially fit the prevailing norm at one time and not at another, or not in another crowd of people. (Compare, "certain English sentences today were not English sentences in 1400, and are not Pidgin English either.")

3) Anthropology, I think, must also have a lot of its componential analyses on a unique basis. For example, the study of culture often exaggerates the separation (in primitive society) between various aspects of the society: the "social" is overly differentiated from the "economic" and "religious," and the true lines of relationship-- the functions of trival societies, artifacts, etc.-- are obscured in the text, as they were not in the field work, by arbitrary divisions. I make here what resembles a Gestalt-claim: that you cannot break down trival structure into "religious," x "economic," etc.; as components, but just as the woof of a rubrick that the significant components will be culturally unique.

4) The study of language and style: Contemporary linguistics has made the study of language the study of the formulae for the formation of correct utterances. The formula is a paradigm for

for the systematic, which is the range of possible sentences. The study of style, in anthropology and aesthetics, also seeks to induce formulae for sensible description in more than impressionistic terms. I feel confident that the style of, say, van Gogh, could be expressed close to completely in a formula relating hue, brightness, texture, object size, canvas size, ranges and types of gradation, etc. Each of these characteristics we have induced would then be a feature of the paradigm (which is the formula,) or the systematic (which is all the occurrences of van Gogh paintings.) There may be confusion here about the formula. To say that a formula exists is not to say that something is not brilliant, inspired, beautiful; neither is it to say that once the formula is approximated, a machine can turn out a new van Gogh. The formula is rather a set of inductive facts about a series of things, comprehending the facets they have in common. (But we said, "close-to-completely.") In any case, the van Gogh systematic is all the paintings he did, and the allodigms are the major types of painting and stylistic periods.

One interesting thing about language and styles, to me, is the way a lintuistic formula or stylistic formula can change over time. The change from T_1 to T_2 will be comprehended by one schematic, the change from T_2 to T_3 by a schematic which may not relate to the first schematic at all; the change may be in a new direction, yet one which could not have taken place were it not for the first change.

5) Perception. A systematic of importance is that of learned perception. It will be recalled that a schematic can be thought of as a way of perceiving a thing. In other words, an iceberg, when seen, can be seen as a large thing on top of the water,

[illegible] be seen as something much huger mostly under water; [illegible] we see a fragment. If we see a photograph and then [illegible] is found to be a bad likeness of someone we know, there is a [illegible] sensation of re-alignment as we re-arrange the features to correspond better to the person we know. This knowledge, either of the person, is a part of the perception. When we look at a painting or listen to music, we relate the painting or music, in form or style, to similar things we have seen before, and our perception of the current works will be a part of that. The systematic or learned perception, then, would be a statement of the transformations that perception can undergo through learning. The more complex case, of course, would be that involving the learning of a complex formula, linguistic, stylistic, etc., and the recognition of its instances, and the creation of new instances.

[illegible] of imponderables, irrelevancies and evanescences. The systems that are hardest to deal with, and often those of the greatest importance, are those whose propensities to shift relationships with small changes are greatest; that is, whose important perceptual and relational qualities will change acutely over something seemingly small. It is not that a car will change greatly with a slight turn of the wheel; but on a narrow canyon road, the general tenor of the situation might very considerably if the wheel were improperly turned.

The difficulty with this formulation, of course, is the definition of "small change". The slight turn of the wheel is by inspection a small change. We cannot call it small by virtue of its likelihood, since it is extremely unlikely, that our driver, who has driven this road many times before, will make that small

change that plunges him over. Yet it is still a small change. Statistical possibility is even worse off: just because he has never slipped yet does not mean he never will; just because the human race has never been destroyed by atomic warfare, we have no guarantee that it will not.

Perhaps the small-change idea will work in terms of likelihood after all. But what we must know is the likelihood of different kinds of changes, and how they systematically relate to the environment of the driver: what causes they might have. It is very unlikely that his spark plugs will suddenly foul. It is possible that he might go over the cliff either because of a hole in the road or because of his psychological state.

I have dwelt on this example because I want to say that there is a number of things of importance for study whose general contour changes markedly with small changes: these are the mystic's ineffables and imponderables, like "experience," "feeling," "culture," "mazeway" (anthropologist Anthony Wallace,) and so forth. These vary enormously with happenstance. It is, naturally, this enormous variation which has inculcated in psychologists such a deep and heartfelt repugnance for the analysis of phenomenology. Indeed, this care to shy the treacherous has evoked many felicitous discoveries. However, I cannot understand the point of view that denies that subjective experience will never have any place in a science of psychology. To justify this, a psychologist would have to hold one of the following viewpoints:

Subjective experience is

a) Existent but in no way discernible

b) Existent but <u>hard</u> to discern

c) Existent but immeasurably, therefore unworthy

d) Possibly existent, but theoretically unnecessary

e) Nonexistent.

I think that only the second claim can be substantiated; even in reply to that, I am certain that a few formulae for necessary subjective relations may be developed. In any case, I am certain that we cannot call phenomenology unworthy of study, since "worthy of study" generally means "someone wants to study it."

PART THREE. NORMATICS.

I do not have space adequately to deal with the normatic, or systematic of choice; I will confine myself to a few definitions and the statement of the normatic principle, without remarking on the utilitarian economics of the business.

The normatic is the systematic of choice and planning. As such, it has a schematic base; and the schematic has three kinds of features. 1) The focal feature, or "what we are aiming at." 2) The instrumental feature, or that which has to be done anyway to maintain life, make it possible to achieve the focal feature, etc. 3) The incidental feature, which is any other desideratum, along the way: freedom of speech, good food, etc. Now, the range of the normatic/focal paradigm is all the normatics that will achieve the focal feature. Since this excludes any wild plans which would omit instrumental features-- that is, can not get off the ground-- the choice among normatics must be a game-theory choice according to degree of fulfillment of focal and incidental features, based on a weighting of which ones are most desirable. The incidental

The fundamental normatic question is this: what schematic of the problem are we to accept as our paradigm? This will have strong consequences for the solution of the problem; because the way you define the desiderata, and consequently the measures you take (and what things are lumped together,) will shape the future beyond your planning (perhaps.) In other words, if the desideratum is "eliminating crime," then we are likely to treat all crime in generally the same way. However, I do not preach General Semantics: it is not the way we talk about the problem.

but the way we treat it, that counts. To phrase it differently: "When can we allow exceptions?" The answer is, "When this case is different." But we want some objective idea of "different." If we seek the greatest utilitarian profit, we should make an exception when the exception brings us (gambler, society, etc.) more than the long-run consistent behavior.

Should all criminals be paroled? No, if there is an isolable group of whom we can predict recidivism with high accuracy.

Should we subject everyone to the same introductory courses? No, if we know that a selectable group can learn better in another way. Etc. There is a compound effect: against the benefit of allowing exceptions, we must weigh 1) the effortlessness of the status quo compared to the cost of administering the new system, 2) the cost of the change-over, 3) the cost of estimating these costs and designing the changesover.

The definition of the job also has a powerful effect. When rules are made, committees formed and offices put in charge, it is done so according to a schematic which corresponds in some way to the features of the normatic involved: committees will be constituted to handle specific problems; bureaux will be compartmentalized according to the general headings thought to be most summary and effective.

Example: in a theatrical production, the central hierarchy of control lies in the producer, the stage manager and the director Each has an integrated task to perform: producer gets the money and sees that it is well spent, stage manager seesk to technical detail-work, and director unifies the production as anm aesthetic object. This is not a necessary structure, it is empirically the best one for general purposes. The unity of each job will collapse

under special strain and extraordinary circumstances, and it would be found that other schematizations would have more salubrious effects.

Even our "best normal structure" sometimes gets out of hand: one of the three can be incompetent, or defect and stop doing his job properly. If it is the stage manager, all is lost. If it is the producer, he may be persuaded or bullyragged by the others to accept their judgements for the good of the production. If it is the director, the actors may sometimes be able t- if they are extraordinarily resourceful-- to do a good job without the director's actual help. Of course, it may go to pieces anyway; but these are some prominent ways it could be saved in these eventualities. There is, we see, a particular systematic of the kinds of trouble that can result from the organizational setup: the bumbling or defection of any of the three will mess up his particular sector of control.

I would like to make this a general principle for all normatics: the sectoring of control and jurisdiction, based on the schematic of normatic features, will have a particular systematic of smooth functioning and a particular systematic of disruption. Whether a setup is the <u>best</u> setup for a given problem or for sustaining copage with something will depend, first, on whether it is generally effective; secondly, on the pattern and possibility of things going wrong; thirdly the alternative normatics which would have contrasting setups and hence contrasting results.

The reason we prefer democracy to dictatorship is not the desire of the people for participation, but the empirical presupposition that it is safer than a dictator. We know well that an autarch is better capable of comprehending a far-reaching and many-

sided situation than a representative assembly; what we do not trust is his likelihood of perpetual fidelity to the ultimates and proximates of popular desire.

Similarly, the way the Cabinet Posts are selected contains implicitly the likelihoods of kinds of solution that can be offered for a given problem. If we have secretaries of Labor, Commerce, Transportation and Health, the kinds of action and policy they will want and be able to effect will be different from that of a Cabinet having Ministers of Youth, National Morale, Entertainment, City Planning and Equality.

Appendix B
Walkthrough of the 1958 Schematics Paper

phSKX58walkthru
01.11.01 (d9

To Ted Nelson Home Page

WALKTHROUGH OF THE 1958 SCHEMATICS PAPER,

"Schematics, Systematics, Normatics" by Ted Nelson

Written by him some 43 years after the original paper.

This is a loose page-by-page explanation of some of the things that were actually *said* in "Schematics, Systematics, Normatics", as distinct from what was *meant*, which is discussed elsewhere (main body of the text of T. Nelson, *Philosophy of Hypertext*).

I will only improve these notes in the unlikely event that someone wants to study this paper further for its historical interest or unusual scope.

(The following page numbers are keyed to a division of each original page into five horizontal strips, a through e. Thus p.10a is the top of p.10, p.11c-12b goes from the middle of p.11 to just before the middle of p.12. A short page, such as p.30, has fewer strips, each being 1/5 of the page.)

P.1: STRUCTURAL MODELS

This introduces the notion of a description or structural model ("schematic"), expressed as a set of relations among elements. Commonalities among such models may be recognized as similarities among the elements and their relations.

P.2
The application or correspondence of a schematic to the world is an interesting property, rather different from the truth or falsity of a proposition.

Transformations among schematics may be interesting.

Parts of a schematic (dyadic relations etc.) may be considered.

p.3-4
The notion of parsing, or taking apart, a schematic as a system of relations (awkward God example)

p.4
Measuring the complexity of a structure of assertions is not easy.

P.5 Shapes of schematics--
"Circular schematic" refers to two things--

- a structure of elements and relations which might be represented by a circular diagram
- a consistent system of assertions, which has redundancies of various kinds

Enumerating the different "shapes" is meant to refer to the structures of different models.

p.5e
"The holistic schematic does not exist"--

This meant that "there is no such thing as a template which is intrinsically unified." One system of beliefs or assertions is not necessarily more unified than another, although someone who holds it may think so.

p.6 [exceptionally garbled]
"Dynamic schematic" was meant to be several things--

- a set of dynamic relations, e.g. a model of ongoing interactions, feedback etc.
- The postulation of "dynamic entities" in a social-science model of the Freudian type, which postulates unseeable relations between unseen entities.

p.6g 'Coherence.' Let me restate this as follows.

The inconsistency of a human set of ideas is not necessarily easy to establish. The believer can always take various steps to shore up a position. People who believes in conventional logic do not recognize that such modes of thought, taken to salvage a position, can be a legitimate form of cognitive endeavor in many areas.

p.6-7
"Schematic" is here propounded to mean *logical system*. It can be expressed either as a diagram or as a set of propositions.

P.8-9
An asterisk notation is proposed for expressing complex sets of relations. If two elements are related, call it a*b.

LET'S BEGIN WITH A SLIGHTLY CLEARER NOTATION--
First consider a parenthetic notation, which would be simpler than what we see here: If a relation exists between a third element c and a*b, call it c*(a*b).

NOW, THE NOTATION AS PRESENTED--
the use of multiple asterisks, was intended to express these matters without parentheses. Unfortunately it makes it more difficult to grasp.

a**b*c was supposed to mean a*(b*c), "a has a relation to the related pair of b and c"
and so on.

EXAMPLE. Consider a construction of strings and sticks. Two strings hanging at both ends of a horizontal stick would be a*b. Hanging this horizontal stick from a string would be c*(a*b).

This proposed, implicitly,

- a linear notation for expressing structures
- a linear notation for constraints
- a generalized expression of all structures in terms of constraints
- a Lisp-like system of expression for constraints and general structures.

p.9-10

"How many relationships are there?" Clearly, from a structure of relationships can be generated an infinity of compound statements expanding these assertions.

p.10c
Anything which matches the structure topologically is the same structure.

p.10e. 'Rules of transformation'--
Here I was calling for both a unifying diagrammatic way to represent structures of ideas, and some way of transforming the diagram or structure to make other assertions in the complex more visible.

p.10-11
"Order of a schematic" in which "relations-to-relations continue to be important"--

I believe this referred to what specifications need to be maintained and which ones become dont-cares, expressed relationally.

Example. I want to keep track roughly of the amount of change in my pocket, in order to give exact change for small purchases, but not the exact number of each type of coin.

Example. In the modern American kinship system, we track relatively few different types of relatives, and we lack terms for many of the relations which in more traditional societies would be carefully noted. Thus a son of my brother is my nephew, but there is no particular name for my nephew's wife (a dont-care, or relationship to which the system is not extended).

p.12 'sameness' as logical identity
Everything that matches a structure is an instance of it

p.13e
'We can see anything in the world in terms of any schematic.' This is not merely a statement about human stubbornness, but also about the plasticity of ideas.

p.14a-c
Similarity is finding arbitrary resemblances anywhere in the web of complex relations that constitute an object. SUCH RESEMBLANCES

COULD BE EXPRESSED IN THE PROPOSED ASTERISK NOTATION.

Extrapolating this idea as intended, this meant that any resemblance -- and THEREFORE ANY GENERALITY-- could be expressed as a set of relations written in some such structural notation.

(EXPANDING WHAT IS IN THE PAPER: The resemblance or generality thus found is an expression of relations which were present in each instance, found to be had in common by the different instances-- two instances in the case of a resemblance, many instances in the case of an empirical generalization. And so in principle, if we could enumerate all the relations which exist (so often an infinity, of course), we could exhaustively compare these specific one-by-one enumerations to find all resemblances and generalizations. In a very very few cases this can be done.)

p.14e
However, we tendencies to notice and distinguish particular things and particular differences.

p.14e-15a "Paradigm"
(Note that the word "paradigm" had not yet been given assigned its present meaning by Thomas Kuhn (1962, some four years later).)

I am using it here to mean a structural model to which something else is going to be compared-- a reference structure.

p.15
Enumerating features of the model

p.15b
By "taxonomic notation" I mean what is now called "lexical order".

This difficult passage is about comparing instances of structures ("schematics") to a reference model ("paradigm").

"Paradigm series" apparently means all the diffferent structure instances which are being compared to the reference model.

EXAMPLES:

•• a description which is being applied to a lot of different instances (e.g. "comedy")
•• a rule which is being applied to a lot of different cases

p.15c
Having insisted that structures are unified, we now break them down into "features", or component relations with letter names.

p.15d
This difficult passage is about mismatch between a model and the instances it is applied to.

• "In each case we can say the schematic fills the consecutive features with some leftover (entities and relations not part of the paradigm)..." This is attempting to say: For each feature, we may find some ENUMERABLE discrepancy or mismatch between the instance and the reference structure.

p.15e
We may be able to find some useful quantitative difference between two structural models.

p.16c-e "Allodigms"
By this I meant: what matches a model can have various structural variations.

Examples:

• Ford cars come in several types
• The same word is recognizable in different dialects

p.17c-e
We look at some instances of meaning.

• such terms as vicious circle, tight squeeze, victim, keystone, melting pot may clearly be relationally defined, as may combinations of them (as in legal propositions)
• however, 'the universe is an elephant' is not neatly defined as the combination of relational structures.

p.18a
It can be a mistake to try to make things exact too soon.

p.18b-c
In many contexts, precision is not necessary. In the warning "Watch it!", the person being warned can probably figure out what the problem is. In the example of the coffeepot simile, the correspondences to 'life' can presumably be recognized by the hearer.

p.18c
The segmentation issue: how a continuous structure is divided into components.

p.19a
By 'systematics' I will mean the comparison of real objects to templates [called 'paradigms' in the paper].

p.19b
A systematic is a range of variations of a 'thing' (structure or objects constituting a class).

p.19c
'Things' do not exist, but are rather structures and patterns distinguished by our perceptions from amongst a much larger flux. The grouping and boundaries of a class are more arbitrary than they appear.

A 'schematic', then, is a perceived or postulated structure used as a template against the world, and to delineate a class.

The 'systematic' of the class is its range of variations [and rules of co-variation-- how varying traits are correlated].

p.19d
Cow example: the perceptual and conceptual template for the concept of a cow. This includes central features (such as horns and udders) and correlated features (such as disposition).

p.19e
[The intended meaning]
A 'systematic' consists of all the traits, and systematic variations of traits,

of a class.

p.20 OVERALL
Example: FLYING DISC ('frisby') trajectories.
This page postuates that there is there might be a dependable similarity formula for the flying-disc throws of every individual, and indeed for sets of people genetically related. This might be expressed by a relational formula-- a flutter pattern, a pattern of lift and downturn; and these elements might in turn have other correlations.

p.20a-e
The flying-disk throws of "B.P." (Barbara J. Price), whether long or short, have a characteristic flutter in the late stage of flight, which sometimes smooths out.

This set of similarities suggests that we could find a characteristic formula for her flying disc throws.
Whereas the flying disc throws of brothers "D.S." and "P.S." (David Schickele and Peter Schickele) are extremely similar to one another, having a characteristic flutter during about the first 1/6 of the trajectory. This suggests that some inductive formula could correctly describe what the flying-disc throws of the two brothers have in common.

p.20e footnote
This semi-humorous footnote suggests the question of whether a genetic basis exists for this similarity (which would seem probable) and whether it could be localized in a specific gene (which seems unlikely).

p.21a
Let us extend the concept that every individual has a characteristic flying-disc throw ('frisby thrust'), which could be expressed by a relational formula; whose elements might in turn have other correlations, systematic patterns of variation.

We may then hypothesize a generalized master formula which would compress and contain the other examples, (much as the general model of Germanic languages compresses and contains German, Dutch, English and Norwegian).

p.21, para 1 last sentence
('Factor analysis' here is used in a more general sense than the usual

statistical method.)

p.21 para 2
The phrase "variations-from/instances-of a central pattern" summarizes the idea of a 'systematic' here.

WHAT THIS MEANS: Given a central pattern or template, we may express, recognize, or find deep commonalities in terms of such a central pattern or template.

p.21b-c
Let us take this idea to some of the issues of the social sciences, especially the study of attitudes, values and beliefs.
Such studies generally divide attitudes, values and beliefs along highly arbitrary dimensions. [Note: these are the concepts embodied in the Allport-Vernon Values Test and Frenkel-Brunswick's study of the Authoritarian Personality.]

p.21d
Instead, we should go about it by "the systematic method, whereby the persistent relations are uncovered".

[By this I believe I meant an exhaustive enumeration of co-occurrences across a vast canvas of instances, and codifying them into a formula stating what combinations do and do not occur-- very like B.L. Whorf's enumeration of the English monosyllable.] [The author later refers to this as the 'factorial' method, though it is not factorial analysis in the usual statistical sense.]

Such a method should uncover "major patterns of variations, minor variations, etc." and uncover "major contours of belief".

p.21e
The resulting findings would better correspond to the actual viewpoints and argumentation of people's thinking.
It might be argued that this is what such studies did already, discovering these lines of division into six different dimensions [the Allport-Vernon value model], but I question this--

p.22
Conjecture: the 'systematic' approach [enumerating and codifying

correlations across a vast canvas of instances] would uncover "major contours of belief"

22b
By such a method, we would uncover patterns of variation that would not map well to any template ("imported schematic"). Such difficult-to-map variations we may call "idionomic". [The term "idionomic" is here imported from psychology, where it is distinguished from "nomethetic"-- meaning mapping to some generality.]
The conventional method is "plumbing inaccurately the same thing we wish to sound systematically"-- meaning that the proposed enumerative method should have much greater exactitude.

p.22b-c
Some things may not the given template-- the prevailing norms of thinking, or the English language. (The author uses the term "factorially fit", meaning "fall within the enumerated scope of systematic variation".

22d
"I make here what resembles a Gestalt-claim" means the author believes that breaking something down into these parts leaves out some important aspect of the whole.

Specifically, that tribal structure does not have parts like "religious" and "economic" as components [separable entities] but as "the woof of a rubric"-- that is, analytic distinctions of parts that may be enumerated but are actually inseparable.

p.22e
Linguistics is now "the study of the formulas for the formation of correct utterances." Such a formula is a template ("paradigm")--

p.23a
for the range of possible sentences.

p.23a-b
Let us extend this model to finding formulas for describing style in anthropology and aesthetics. Example: we might be able to express the style of van Gogh could perhaps be expressed "close to completely" in a formula expressing certain properties. Each of these properties would be a 'feature' [parameter] of such a formula.

p.23c
To describe a work of art formulaically does not deny its brilliance or beauty, and does not assure that a comparable work could be created automatically from such a formula.

The term 'allodigm' (introduced earlier) is here used to mean one systematic variant in a model or template, by analogy with "allophone" (linguistics) [and now "allele", in genetics].

p.23d
Such formulas-- e.g formulas for language or style-- may change over time and in unexpected directions which do not clearly derive from one another. For instance, if T1 is a first formula (for instance, for a grammar at a given time), it may change to T2 (the grammar two centuries later) by one set of transitions, but then T2 may change to formula T3 (for instance, the grammar two more centuries later) by changes seemingly unrelated to the first set of changes.

p.23e
Learned perception is a systematic, i.e a set of variations of formulaic relations.

p.23e-24a
Iceberg example: First we perceive an iceberg as a large object floating above the water; later we learn that this visible part is the tip of a much larger object floating UNDER the water.

p.24b-c
As we learn more about a class of objects, their variations become apparent. We increasingly recognize their instances and transformations.

p.24c
'Impalpables, ineffables and evanescences' are often those perceived structures and objects which have complex formulaic structures, but change greatly as variants within those formulaic structures.

p.24e
However, this issue is made more complex by the problem of saying just what is a small change, since we have no formulaic way to recognize size of change.

p.25c
Many 'ineffable and imponderable' structures-- feelings, culture, etc.-- can change enormously with small changes.

p25d-26b
Psychology has attempted to expunge the analysis of feeling ('phenomenology') from the field, but this cannot be justified.

p.27 OVERALL
(This is an attempt to state some of the problems of utilitarianism, the seeking of the general good. Utilitarianism is generally formulated as some simplistic decision problem, like "What will bring the greatest benefit?" However, real-life decisions about benefit require translating the benefit-calculus into some system of rules and generalities. These remarks about normatics are a condensed, telegraphic way of trying to express some of the problems of creating rules, generalities and moralisms to support any particular utilitiarian calculus. They are further generalized around the problem of formulating rules, generalities and moralisms to support any system of desiderata, whether traditional or religious or not.)

'Normatics' is the application of these concepts (STRUCTURAL MODELS AND FORMULATIONS OF VARIANTS) to choice and planning [and later, implicitly, to ethics and morals].
Most broadly stated: By our rules and moralisms we seek to achieve certain goals in the world, or maintain a way of life. Those rules and moralisms may or may not contribute to this objective.

p.27b
A normatic structure ('the normatic') is based on a schematic (template).

[The following is not part of the central argument.]

A normatic structure has three aspects-- what is intended; what act or rules ('instrumental feature') are chosen to implement that intention; and what other things need to be part of the plan ('incidental feature').

p.27c
The 'range' of the normatic-- that is, the variants of the template-- includes

the different variant rules that express it, even if some of these rules cannot work usefully.

p.27c-d
Choosing a set of rules or moralisms [assuming a utilitarian point of view] has benefits which are hard to project in practice, and thus bring in various game-theoretical issues.

p.27d-e
"The fundamental normatic question is this: what schematic of the problem are we to accept as our paradigm?" This means--

- depending on how you formulate the problem or the desiderata, you will come out with different sets of rules and moralisms
- what desiderata and categories ('things') you seek to optimize, and recognize in your rules and moralisms, strongly determine the outcome

p.27e
Example: if we consider 'crime' to be a single category, then we look for solutions that deal with all crimes lumped together [whereas if we distinguish between major and petty crimes, violent and white-collar crimes, it could lead to more nuanced proposals and solutions].

p.27e-28a
However, it is not merely the concepts by which we refer to problems and desiderata, but what we do on the basis of these concepts, that is important.

p.27a-c
For instance, the question of when to make or allow exceptions. If the exceptions to be made are easy to recognize and have well-defined consequences, this justifies a system with exceptions.

p.28d
Rules, committees and departments are examples of normatic structures. The division into names and headings comes from the conceptual structure and headings ('schematic') with which we approach a problem and which we believe will be effective.

p.28e-29a
Example: division of theatrical roles into producer, director and stage manager. This has been shown to work, but can fall apart under special circumstances.

p.29b
If one of these three runs into trouble, other participants may be able to rescue the situation.

p.29c
This division into roles can engender characteristic sets of problems.

p.29d
'the sectoring of control and jurisdiction ... will have a particular systematic of smooth functioning and a particular systematic of disruption.' This means that a standard setup will have standard ways of functioning well and standard ways of going wrong.

p.29d-e
To choose the best setup, we consider alternatives. We must look at both the patterns of effectiveness of the choices, and their dangers.

p.29e-f
Dictatorship is in some ways a good system-- the dictator can consider more aspects at once-- but dictatorships are likely to go very wrong from the standpoint of everyone else.

p.29b
The division of functions for cabinet ministers will have different characteristic approaches depending on how the cabinet posts divide up the universe.

•

•

Appendix C
Hypertext Notes, 1967

WORKSPACES, TEXT THEATERS, AND HYPERTEXTS: PURPOSE OF THESE NOTES
Theodor H. Nelson

Hypertext Note 0 [19 Apr 1967]
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

Interest in hypertexts leads directly to a concern with text theaters. You cannot make or present a hypertext without a text theater; and the less limited you want the hypertext to be, or the less sure you are of what your final product will look like, the more versatility you need in your theater.

Likewise, unless one wants his hypertexts to conform to fixed and simple rules, and appear in fixed and simple formats and activity spaces, then the nature and creation of activity spaces becomes of great concern.

Thus there are several things one can do about hypertexts.

1) One can try to create hypertexts without a theater, which is difficult and probably pointless.

2) One can endeavor to figure out in the abstract what kinds of hypertext (and activity space) are possible, and their properties.

3) One can design text theaters in which the creation and use of hypertexts (and activity spaces) is possible and practical.

These hypertext notes report principally on work of types 2 and 3. Because of the interest in this work among those in computer text handling, these notes are intended to short-circuit the publication cycle, and get various ideas abroad expeditiously.

BRIEF WORDS ON THE HYPERTEXT
Theodor H. Nelson

Hypertext note #1. 23 Jan 67
Private circulation only. This will eventually be incorporated in a longer article.

About the term. 'Hypertext' is a recent coinage. 'Hyper-' is used in the mathematical sense of extension and generality (as in 'hyperspace,' 'hypercube') rather than the medical sense of 'excessive' ('hyperactivity'). There is no implication about size-- a hypertext could contain only 500 words or so. 'Hyper-' refers to structure and not size.

What kinds of structure are possible in a hypertext? Any. Ordinary text may be regarded as a special case-- the simple and familiar case-- of hypertext, just as three-dimensional space and the ordinary cube are the simple and familiar special cases of hyperspace and the hypercube.

(It should be understood that all other examples of hypertext will also be special cases. There is no 'true' form of hypertext, only a very large number of possible structures.)

Ordinary text, called by McLuhan and others "linear," is a continuing sequence in a fixed order. (Actually there are many departures from this: the footnote is the most obvious, but the summary, illustration, headline and other features also break away.)

But from the standpoints of both subject matter and presentation, it is rarely ideal to have a single fixed sequence of materials or ideas. And there are ways now to present useful text structures of every conceivable kind.

To generalize the idea of ordinary text, we need to find a means of representing departures from mere sequential order. This lies in the mathematical theory of graphs. A graph structure is a web of connections between points. It may have connectors that have directions ("arcs")

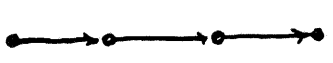

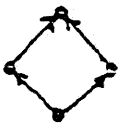

or it may have undirected connectors ("edges")

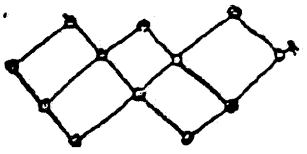

and it may even have arcs or edges of different kinds ("colors"):

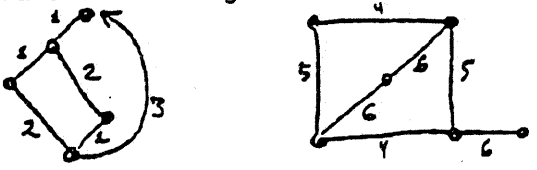

Arcs, edges and colored connectors all have possible analogues in the hypertext.

Each dot represents a chunk of text. The lines represent connections between chunks of text. A plain line, or edge, represents a connection of possible travel from one chunk to another. A directional line, or arc, represents a connection of possible travel from one chunk to another in <u>one direction only</u>.

Ordinary text is like this:

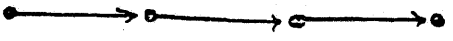

However, since it is possible to backtrack or skip, it is really more like this:

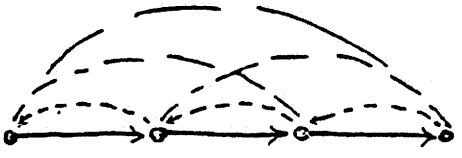

In this case, the arcs represent prefabricated or convenient connections (directional), and the dotted lines represent optional travel (in either direction).

Finally, "colored" arcs and edges-- that is, arcs and edges that have been assigned to different types-- may represent <u>conditional</u> travel. For instance, the author/editor of a hypertext might arrange that only persons having certain qualifications (or persons who had covered such-and-such material) should be allowed to pass along a certain connection. The numbered (or 'colored') connector types may represent these permissions, and other conditionals.

Hypertext diverges from ordinary text in that the reader's possible sequences diverge from plain sequence. Is he to have choices? And how are they to be expressed?

The answer is that the reader has the choices that the author/editor allows him. He may not be aware that choices are made, and may think he is reading a book and expressing his opinion of it. Or at the other extreme, he could be given complete and continuing information on where he is in the overall text.

In ordinary writing there are two problems. The first is structure: creating a suitable and orderly division of the subject, and creating appropriate sequences and sections in the text for its proper presentation. The second problem of ordinary writing is presentation: keeping the reader aware of what he has already read, keeping him oriented as he sees different parts of the subject, and keeping him interested.

These problems are unchanged in the hypertext.

Hypertexts are inevitable, and will come into being just as fast as text handling systems come into the field. Several systems suited to handling some hypertexts exist already: Magnuski's DOC system at Project MAC, Bernstein and Slojkowski's program documentation system at Bell Labs, and Engelbart's editorial console at Stamford Research Institute.

HYPERTEXTS and, or perhaps versus, CAI
Theodor H. Nelson

Hypertext note #2 23 Jan 67
Private circulation only. This will eventually be incorporated in a longer article.

Programmed teaching has been improved by the introduction of computers. These make it possible for a teaching program to branch, and follow a graph structure. (Thus it may, of course, be a type of hypertext.)

Practitioners of this art-- the stringing of program "frames" for teaching-- are generally of the behavioristic persuasion, believing a) that it makes sense to define the goals of teaching simply in terms of the subject's resulting behavior, and b) that a construal of the teaching task in terms of reinforcement schedules and evoked behavior successfully comprehends the whole problem. To question these assumptions is pointless, since they are either self-evidently false or there is no point in arguing.

Anyway, it is believed in this field that since the problem is one of "learning" (in the pigeon/rat sense), the proper way to teach is by a highly detailed succession of presentations and fill-ins, minutely crafted to provide explanation and "reinforcement" in precise and tiny doses.

The alternative viewpoint I would like to offer is that the problem of teaching is motivation (and comprehension, which raises or lowers motivation) and the way to teach is to present things so as to maintain high interest. It is necessary to permit individual, personal success in an individualized study environment-- where the student makes a field his own, immersing himself in those aspects of the field that appeal to him, probably without the interfering personification (and perhaps gratuitous interference) of the teacher. Understanding, autonomy, success at exercises, and motivation must all grow together.

Large hypertexts, with many divergent but cross-cutting regions of study and activity, seem to be the way to do this.

Without harping, let me list several salient distinctions between this hypertext approach to education and the CAI approach.

1. Reinforcement vs. Motivation. A motivated individual will arrange to be reinforced. A reinforced individual will not necessarily arrange to be motivated.

2. Minutiae vs. sweep. CAI requires infinitesimal attention to particulars. Hypertexts require considerable attention to generalities.

3. Memorization vs. Comprehension. The old sense of comprehend was "comprise"-- press together a variety to form a whole. (This is the whole point of survey courses and outside reading; yet by making these compulsory the most important factor, selective personalization, is lost.) Memorizing is not understanding.

4. Acceptance vs. Considering Alternative Viewpoints. All fields have areas of controversy, or at least (as in mathematics) entirely different ways they can be considered. Rather than enforce the initial acceptance of a particular viewpoint-- which is why children disaffect from subjects and from school-- we must give variety and alternatives. What we like best we learn best; and this is no dead end, but a framework for the rest.

5. "Learning the basics" vs. Giving Alternative Handlings and Dismantlings. The myth of education is that there are "basics" in every field. This is simply not true. There are indefinitely many ways that each so-called subject can be dismantled for teaching. Good teaching provides handles and approaches for a sense of understanding, and alternative sets of "basics" in the hope that one will appeal.

6. Difficulties of Production: Crossword Puzzle vs. The Panorama. The difficulty of making good teaching programs is notorious. As a form of activity it resembles nothing more than creating crossword puzzles. Will it be like this to write teaching hypertexts? Probably these will be difficult, but not so compacted and intricate. The central thing we may depend on is this: someone who knows what he is reading or doing, and has freedom of movement, can be relied on to get himself out of a meaningless dead end. If he is being led blindfolded, without explanation, through a maze of responses and reinforcements, everything must be correct; for if he is stalled he cannot help himself, and he cannot help you.

Another possible analogy is the difference between making animated cartoons and epic films. They both take money and trouble. One requires great attention to specifics of detail in interlocking fragments; the other requires equipment and material on a grand scale, and logistic and aesthetic direction.

We may think of CAI as an extension of programmed teaching, of the general hypertext as a generalization of writing and editing. Whether findings from CAI will have a bearing upon hypertexts for teaching may be doubted.

A word on Socratic CAI. "Socratic" systems have a touch of humanism in them, as they try to express conversationally the exact understandings and misunderstandings of the pupil. This is a wonderful technique but is it universal in application? Consider: would you rather have directions to get to a place, or a road map? The former introduces spurious social interaction, an unnecessary problem in most cases, and does not provide for errors either by the director or the follower. A road map allows self-correction.

Moreover, the Socratic system has a tendency, at least in the published examples, subtly to put down the student while it promulgates the myth of the single right answer. This is what we are doing in education already, to the national detriment.

TEXT AND TEXT-RELATED DISPLAY: SOME TERMS

Theodor H. Nelson

Hypertext Note 3 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

There will be many ways of displaying texts on computer display screens, and many reasons for doing so. There will therefore be different kinds of theaters, or settings, for handling texts upon a screen; and many kinds of parts and arrangements in which the texts will appear. In this and later notes various types of setting, and different parts of displays, will be distinguished. I hope these terms will make it easier to describe, and think about, present and future systems. The overall intent is to clarify a lot of needed ideas, and assign the most appropriate terms to them.

Theater. Let us use the term "theater" (or perhaps "setting") for a display console that has various capabilities for displaying, filing and computing. Display system has been in common use, but this term has problems. "System" could refer to the console alone, the program and file facilities available through it, or the overall computer environment-- time-sharing, monitored or whatever-- in which it functions. It is simpler to speak of a "theater," restricting that term to the console and the facilities available through it. By this terminology, then, one system can have many theaters -- text theaters, engineering design theaters, and so on.

After all, the dictionary sense of "theater" is "a place where certain things occur." (Webster III, definition 4a: "a place or sphere that is the scene of dramatic events or significant action.") (I think this term "theater" is also preferable to "facility." "Facility" best refers to some acting part that does a specific kind of job; for which the term "capability" is not really satisfactory, and for which the term "capacity" should likewise not be used-- the capacity of the system being how much it can hold, in its various buffers, buckets, cisterns and oceans.)

A facility, then, is an acting part of the system or theater that does a specific job. Possible text facilities include the filing and retrieval of texts, word and phrase searching, graph following and graph structure analysis, the placing of markers and connectors in the text, and, of course, simple text editing.

An information base is the information on which displays are founded. Modifying the information base may take place in some systems, not others. Text editing is one way to modify an information base; rebuilding a SKETCHPAD structure is another. Different systems will have different facilities for modifying the information base.

The file structure of a system is the way in which information is divided, connected and stored. The extent to which a file structure will influence the material stored on the system will vary. Many features will in part determine how much the user is affected by the file structure. These will include types of connectors, record length, restrictions on connectors, categorization methods and many other aspects.

WORK THEATERS, WORKSPACES AND THEIR PARTS
Theodor H. Nelson

Hypertext Note 4 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

A spread (screenful? presentation? display? frame? shot? event?) is a collection of text and graphics which is output as a complete screen display.

A window is an area on a screen which "looks into" a connected body of text or other information. However, it does not add to or rearrange material found there. If the text can be made to move through the window ("scrolling"), we may also think of this as the window looking at successive parts of the text file. Thus a placemarker or window pointer (or sashweight?) is a pointer into the text file by which the window's location in that file is maintained.

This "window" idea is useful mainly for free text or lists. For coded information--legal, survey, municipal, bookkeeping-- forms are required. A form is a pattern of words or information on the screen to which is added information obtained from a file. A window that is empty, that is, pointing to an empty area, is an unoccupied window.

A fixed window always occupies the same place on the screen. A movable window may be moved about on the screen by the user (either continuously or in jumps.) A dependent window will take its position from the position of something else on the screen. An optional window may be made to appear or not, at the user's preference. Alternative windows are sets of windows of which only certain windows will appear at a given time.

A layout or format (short for screen layout or screen format) is an arrangement on the screen of different windows, forms, system messages, and choices. (However, system messages and alternatives may not necessarily appear on the screen.) Alternative and conditional formats may be useful.

A sublayout (subformat) is a part of a format (or layout) which may itself have alternatives, be replaced, move around, or appear in different formats.

A free message or loose message is a text message not in a window or form.

A face (door? portal? wall? side?) is a format assigned to a position in an overall workspace.

A workspace or work structure (hyper-workspace? display space? implicit conceptual topology?) is a structure of formats, arranged, in a branching graph, so as to succeed one another upon the screen. A graph of the connections of these formats we may call an inter-format graph. Through these formats the user may choose information and displays, negotiate with and modify the information base, call facilities and enter new information, as required. The construction of suitable workspaces for specific tasks will be a difficult and complicated matter, with personal preferences playing a major part.

Steps from format to format in the workspace we may call jumps, steps or perhaps cuts (as in the movies). Conditional, optional and alternative formats might be desirable at specific vertices.

We might use the term hyper-desk for a console at which all your work can be done, or a significant mix of tasks.

The access structure (or access space) of a workspace is a graph of the ways in which specific information may be reached from different layouts. This may not have the same graph structure as the workspace itself, since the same information may presumably often be reached from different formats.

CREATIVITY SYSTEMS
Theodor H. Nelson

Hypertext Note 5 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

Creativity support system (or creativity system) is a term I wish to reserve for a particular kind of text-handling theater and file system.

To explain this, a word is necessary on the creative process.

Constructive creativity-- the kind that occurs in writing, theoretical thinking and design-- has the following stages. A number of parts, or entities, are seen to hang together under certain ideas.

The entities are then arranged or structured according to the ideas. (For instance, in writing, a number of points are put together under an outline.) Typically, however, these things do not fit well together at first, and various combinations and arrangements have to be tried out and compared. Ordinarily we can only try a few arrangements, often only one, which is patched, rearranged and repaired until we have to stop. The ideal creativity support system, then, would help us put together, speedily and easily, all the different arrangements we wished to compare, and let us compare them quickly and conveniently.

The following, then, are the criteria of a creativity support system:

1) The system can hold alternate structures of the same materials, and accept new ones easily. If it is a text system, these materials will be texts; if a system for mechanical design, they might be alternative designs of a machine. (Alternative criterion.)

2) The user may easily compare the corresponding parts of these alternative structures. (Intercomparison criterion.)

3) The user may easily label and annotate these alternatives, in various ways and "levels," so that he can summarize and keep track of his increasing insights about them. (Annotation criterion.)

4) The file categories, label types, etc. must be easily changeable to what the user wants. (Protean criterion.)

5) It must be possible to accommodate entries of any length, or at least very long entries and very short ones.

We are now ready to discuss what I will call the general creativity console.

Such a console will be a creativity system, as described above. Moreover, it must be "general" in that it will lend itself not merely to text writing, but to musical composition, graphic arts, motion picture editing and so on. This does not necessarily mean it can show movies or perform musical compositions. But it must be able to hold written information about these things, and connect usefully to equipment that can show them.

This means that it must be readily hooked up to A) receive signals from other equipment, and B) control other equipment in turn. C) Its visual displays must be easily supplementable with new console button arrangements. (In hardware, for speed.)

More importantly, D) its operations must be easily changeable to new ways of working-- by a simple, noncomputerish command structure, perhaps a pictorial one. And finally, E) the record structure by which things are filed must be prevented from affecting the user: whether storage on disk or tape is by ten-character chunks or thousand-character pages, the user must be shielded from being aware of how these relate to the things he is doing. Otherwise the work may be influenced improperly by the distraction of such structures.

PRESENTATIONAL THEATERS FOR INTERACTIVE MEDIA
Theodor H. Nelson

Hypertext Note 6 7 April 67
Private circulation only. This will eventually be incorporated in a longer article.

Computer display makes possible a number of interactive media-- information media showing prepared materials and allowing the user various choices of what he will see. These include the hypertext, hypergram, annotated tableau and hyper-comic.

1. Texts and Hypertexts.

A text-- in the ordinary sense-- is a congeries of written words. These words are arranged in divisions and sequences and connections that seemed to the author or editor appropriate to the subject.

A hypertext, therefore, is similarly a congeries of written words, similarly arranged in divisions and sequences and connections that seem to an author or editor appropriate to the subject. But the hypertext is not restricted to a simple succession, and may therefore branch in different directions. Thus a hypertext is text arranged in a graph structure, with the branches to be made according to the choice of the user or the criteria set down by an author or editor.

The term "hypertext," as used above, can include anything from through-composed text structures to text structures which have merely been collected: as long as they are composed mainly of text sections and branches.

Presumably hypertexts may contain graphics of various kinds, described below. These might be called "illustrated hypertexts," or we might not bother to make the distinction.

2. Graphics.

A "graphic," by convention, is some sort of picture, symbol, design, or other intentional part of a visual display that is not made up of letters. We may distinguish among several kinds of graphics for computer display.

A literal graphic is a graphic which is pulled out of memory and put directly on the screen without variation or movement.

A sketch is a graphic which is modified or translated by the machine, adapting or abstracting from some data, but where the data was created for or in part resembles the final result.

A generated graphic is one which is created by the system from data, mainly ad hoc. The elements of its structure are found rather than prepared beforehand.

Moving or movable graphics are those which are capable of internal movements (as opposed to creeping or relocatable graphics, which simply occupy successive positions).

A hypergram or hypergraph is a graphic, or diagram, that can be manipulated, either freely by the viewer or by the system on the basis of data or other occurrences.

A tableau is a spread which is largely pictorial or graphical, but contains inset text matter in the form of captions, dialog, etc. (If the subjects are people and dialog, "hyper-comics" will perhaps be the better term.)

Such devices may seem far-fetched. However, their use for teaching the undermotivated and the poor reader may be considerable. And their possible use for more serious tasks is not unthinkable. According to one reporter, the President's guidebook to our nuclear warfare command codes is a sort of garish comic book, "pages of close text enlivened by gaudy color cartoons. They looked like comic books-- horror comics, really, because they had been carefully designed so that ... military aides could quickly tell him how many million casualties would result from Retaliation Able, Retaliation Baker, Retaliation Charlie, etc...."*

---

*William Manchester, The Death of a President, quoted in The New York Times, March 29, 1967, p. 40.

These terms may seem unnecessary. But they may help in simplifying the description of existing and possible display systems.

For instance, SKETCHPAD is a graphics drawing facility which permits the modification of a graph-structured data base by screen manipulations. It creates sketches from this data base, and will permit the creation of hypergrams in two and three dimensions.

Or, a picture of a machine or human body which could be "opened" and "manipulated" on the screen, with expandable labels, would be a hypergram tableau.

HYPERTEXT STRUCTURES
Theodor H. Nelson

Hypertext Note 7 26 April 67
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

There are a number of different types of hypertext. They differ principally according to the freedom of the user, whether and to what extent they are hierarchically structured, whether they are continuous or in discrete parts, and the way they are addressed by the user.

Their editorial qualities, and editorial techniques that may be useful, will not be described here.

(It would be convenient if we considered all features of hypertext to be "structural," a matter of graph connections. This is not the case. A number of interesting properties of hypertext are not structural; and hypertexts with 'the same' structure may be presented in ways that are importantly different to the user. The way the system faces the user is rather independent of the connective structure of the text.)

A text is not a hypertext if interaction is irrelevant and the user's actions have no effect on what is seen. If what you see depends on qualities you bring with you, what you already know, or what you are, then it should be called predetermined conditional text, or "fated" text. (However, a hypertext might be partly fated.)

Freedom and explicit control by the user are important attributes. If the user is given orientation and explicit ability to move freely in the text, it is free hypertext. If his movements determine what he sees in ways he cannot control, it is a conditional hypertext, or text maze. If the user is not given overall or orienting information about the text graph he is in, he does not maintain control. Though his actions may determine what he sees, if he doesn't know it and can only guess at the structures with which he is communicating, it cannot be called 'control.' (Conventional computer-assisted instruction is almost exclusively of the text maze type.) A hypertext may, of course, be partly conditional and partly free.)

(A presentation or structure may also be available depending on some action by the user which occurred earlier but not just before the current one. In this case we may call it a back-conditional.)

A hypertext may be hierarchical or non-hierarchical. There is a strong bias among technical people (fostered perhaps by the tradition of the Harvard outline of topics: I.A.1.a. ...) to suppose that hierarchy is natural and universal. This is a complex assumption about the nature of verbal information and ideas. Hierarchical hypertext forms include paragraph trees, stretchtext, etc. But non-hierarchical hypertext forms are possible which could take on any graph structure whatever. How this will usefully relate to content and function, and how the user will keep track of where he is, are problems.

Nevertheless, certain types of non-hierarchical connections should definitely be useful in hypertext. These include "return to start," and material set aside for modal queries (see below).

Continuous and discrete structures are both possible in hypertext, as they are in ordinary text. (By "continuous" we mean something like "smoothly written.") This is partly an issue of whether presentational breaks in the text are formally known to the system (like, in print, the paragraph, chapter break and boldface heading) or simply matters of editorial character.

(The extent to which this continuous/discrete distinction will be part of the computer software and technical properties of hypertexts is not clear. For instance, stretchtext seems intrinsically continuous, but users could elect to give it discontinuities that might be useful.)

The possibility of continuous hypertext with other than a simple sequential structure should not be overlooked. Weird topologies might be useful and possible in some unthought-of way.

The way in which choices and branches are presented to the user may vary independently of the connective structure by which they are strung together. There may be choice points (or breakpoints), where the presentation (or text) stops till the user chooses; or the choices may be imbedded in branch points-- markers in a continuing presentation that indicate the user's branching options. (If we consider that a knothole on a wooden board means a branch once went from there, knothole might be a xylologically appropriate term for branchpoint markers.)

It may be that branch possibilities are not shown at all. (Announced vs. unannounced links.) They might be forbidden to the user, which is probably a poor idea. Or he might have to guess what options are available, by attempting queries. Whether this would serve a function is the reader's guess.

The reader might have the option of suppressing branch-option signals in order to concentrate, and reviving them when he wanted to digress. This would be another matter.

We may distinguish between standing options (always there) and present options (that come and go). There may be modal branch options, different types of branch to connected texts. If all types of modal option are always available, they are standing modal branch options. (Standing modal options could make it unnecessary to signal choicepoint and branchpoint links; there would always be n connections, or perhaps m points of curiosity the user could have satisfied at any time.)

It should be noted that modal branch options are surely not hierarchical, but partake of some other graph structure.

Branches or jumps may occur from any point in text to any other point in text. Ordinary writing usually permits jumps from the end of a sentence or clause to the beginning of a sentence (the footnote); but there are other possibilities.

Branches may occur between two distinct units (like the footnote), or between two places within the same unit, in which case it should be called a jump or hop. (In conventional text they say, "see page 3.") These jumps may be bidirectional or one-way.

(An esoteric question about branching is associated with hierarchical hypertext, such as stretchtext. The question is whether a jump within the same unit should also be to a point at the same level; and whether the branchpoint (e.g. in stretchtext) should be continuously available into further levels down.)

Questions, instead of explicit choices, may be used to manage user response at choicepoints. But this raises interesting issues. Questions may be topologically and functionally identical to free choicepoints. But are they psychologically?

If the questions give no hint as to what the answer will make the system do, then of course their effect is very different, for we have a text maze. But if the user is able to divine the probable result of the answers-- and is thereby motivated to guess, lie, etc., as a way to express choices-- the situation is complicated. Assuming that he is not being judged, graded or punished-- conditions whose psychological effect is not now evident, but it probably won't be good-- there is still that spurious conversational confrontation whose virtues may be less than supposed. (See "conversational systems," another note.)

The connective texture of hypertexts may be of many kinds, following any discrete topology (and perhaps several continuous). The connective texture may follow a regular pattern, like the graphs of mathematical groups; this would be a sufficient condition for creating a standing modal query structure. Or the connective texture of hypertext may be locally variable. The advantages of this have yet to be explored.

There may be one text type (as in most books) or several, modal types, distinguished in their meaning and presentation (like the captions, subheadlines and inset boxes in a magazine.)

Variants of areas or units or sentences may be presented to different users for different reasons; it may not make sense to speak of these as truly alternative units. There is no convenient way to discuss them here.

Randomness of system response might also seem useful, for enlivenment of dull material, challenge, etc. But it might also be demoralizing. The effects of randomness and other de-structuring is not of current interest.

STRETCHTEXT
Theodor H. Nelson

Hypertext Note 8. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article. Terminology is tentative.

Stretchtext is an extremely simple but powerful form of hypertext. It is probably the easiest possible hypertext to understand. It would be hard for a reader of stretchtext to become confused.

Stretchtext consists of ordinary continuous text which can be "stretched" or made longer and more detailed. By pointing at specific areas and pulling the throttle in the "magnify" direction, the reader may obtain a greater detail on a specific subject, or area of the text. The text stretches, becoming longer, with replacement phrases, new details and additional clauses popping into place.

The good of this structure should be evident. The reader remains oriented. If he loses track of where he is, he "shrinks" the text to a higher, smarter level; if he wants to study a topic in more detail, he magnifies it.

An important editorial constraint on stretchtext, then, is that details and narrative arrangements must remain fixed in their relative order through different levels of stretchtext. However, in one respect it appears to be easier to write than ordinary text: rather than deciding what details to "put in" and "leave out," the author merely assigns altitudes (or "fineness"?) to topics and details, thus determining at how great a magnification they will be seen.

High-performance dynamic consoles may be preferable for stretchtext, but ways are possible to present it on a static display-- if the display can support a moving cursor and a user pointing device.

Editorially, the stretchtext is (1) always the same unit, and (2) always a continuous narrative. Thus it is unlike hypertexts with discrete chunks and breaks.

A HYPERTEXT STRUCTURE FOR SELF-TEACHING
Theodor H. Nelson

Hypertext Note 9. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article. Terminology is tentative.

The best teaching is self-teaching. The hypertext design described in this paper is intended to be simple and clear, and help the user stay continuously oriented in a complex body of materials.

This is a hypertext design that can be easily understood-- if not from reading this note, from a few minutes of exploring the text at a CRT.

The design consists of text arranged and interlinked in three different sections; a special user dashboard; and a supporting system of computer and program to make the overall text object respond in the ways suggested.

This design would not be hard to implement right now on a high-performance CRT console with appropriately large memory behind the computer.

These parts are inseparable: the text and links, the dashboard, the computer program. However, the program would presumably be general enough to use with other hypertext designs as well.

The overall purpose of the design is to keep the reader oriented as to where he is and what he is doing. It is intended that the text should be under the full control of the user, that he may quickly find answers to his questions, and find interesting cross-cutting collateral material as well.

As I see it this design is best suited to discursive materials, in particular history or historically-taught survey courses. But in principle anything that can be put in a book can be put into this hypertext design. It has, however, specific editorial properties.

THE TEXT.

The text has three parts.

1) The Narrative, a unit of stretchtext (described elsewhere). At its minimum it is (say) a thousand words. When stretched it gets much larger.

2) The Lookup. This consists of definitions of words (from the dictionary), and stretchtext expositions of certain other terms, and stretchtext biographies. To reach the lookup from any place in the whole hypertext, the user has merely to point to a word, term or name, and press "jump." This definition or exposition then appears on the screen. Because these are stretchtext, if the material is not adequately clear to the reader at the first look, he may magnify it.

3) The Articles-- text pieces ranging from, say, 100 to 10,000 words in length. They may be either ordinary straight text or stretchtext. These contain "collateral material"-- insights, wisecracks, essays-- that tie together specific aspects or ideas also found in the Narrative. These articles may or may not have titles, headings, author credits, etc.

JUMPS.

The text is dotted with jump markers. A jump marker is like an asterisk: a symbol informing the reader that a branch is available at this point. To follow the jump marker,

one has merely to point at the jump marker with light pen or other device, and press the "jump" button.

There are two kinds of jump markers, which we may call N and A. N signifies an optional jump to part of the Narrative; A signifies a jump to an Article. Jump markers of both type N and A may appear anywhere in the hypertext. (The further stretchtext is magnified, the more jump markers appear.) A type N jump marker takes the reader to some place in the Narrative, where the type A takes him to the beginning (or other part) of some Article. (There probably need to be "altitude" restrictions on this that I won't go into here.)

INDICATORS AND CONTROLS.

The dashboard of the system is intended to keep the user aware of where he is, and keep him in control. To keep him aware of where he is, there are three indicator lights, representing the three parts of the text. If he goes from Articles to Narrative, or Narrative to Lookup, he may verify this immediately from the indicator lights.

The user has two simple (but sophisticated) controls. One is a throttle handle which can be moved, by degrees, in four different directions, corresponding to movement in the text: right (forward), left (back), down (magnify), up (shrink). Movements of the throttle effect corresponding movements of the text.

The other control is an Englebart-type mouse-- a box on wheels, whose movements on the table control the movements of a pointer on the screen. There are two buttons on the mouse: "jump" and "return." Various things will happen if you press "jump." If you are pointing at a word, term or name, you will get a definition, exposition or biography. If you are pointing at a jump marker, you will jump to the linked location in the Narrative or Articles.

The RETURN button undoes your jumps, popping them one at a time from a push-down list. Thus no matter how much you jump, it will be fairly easy to keep track of where you are, and get back to where you were.

This hypertext design could be easily made plain to a seven-year-old, would permit great personal option and control, and make possible an extremely rich compound of interesting material in which the user could easily stay oriented.

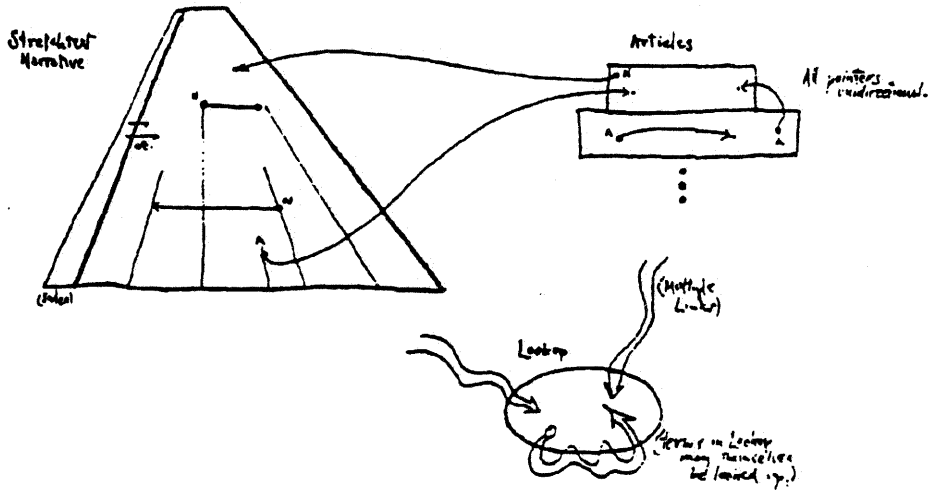

XANADU
Theodor H. Nelson

Hypertext Note 10. 29 Apr 67
Private circulation only. This will eventually be incorporated in a longer article.
Terminology is tentative.

XANADU is to text display systems what SOLOMON was to big computers: an ambitious proposal, dropped by management, which may yet serve to focus thought. Indeed, it may be thought of as exemplifying a class of systems which, unless quite impractical in various technical respects, might turn out to be a prototype for our business, educational, artistic and home text display systems of the future.

XANADU was intended to do text editing; serve as a file handling, information retrieval and management information system; make possible the swift creation of activity spaces for any purpose; and serve for the viewing and creation of hypertexts, as well as lesser text objects.

The general philosophy was that the design of user front ends should be more flexible. In many systems most of the work goes into disk I/O, list processing, etc., and the screen activity design just more or less happens accidentally.

XANADU intended to reverse this, permitting the creation and maintenance of data in any structure, and the ready creation of screen formats and user activity-spaces quickly, at the console, with swift changes possible. The subtlety of screen-work more complex than, say, airline reservations makes it likely that swift change of these activity spaces--'debugging,' upgrading-- will be necessary. Personal idiosyncrasy may also dictate changes and variants in activity spaces. But activity spaces are now only modified at the machine-language level, with quick and simple reconfigurability unheard of. This was the most basic and hardest part. Virtual pushbuttons could be swiftly created, given functions, and then phased out in favor of hardware keyboards. Windows and formats could be quickly made, rearranged

Files. Entries of any length were to have been possible with any number of connector types, and alternative sequences available. (This alternative sequence was to be related to a 'scatter write' facility, by which materials having given formal attributes could be splattered out onto the screen in different ways. (This we might call 'pseudo-sequence.') Storage was to provide for upper and lower case, italics, etc., but ignore them during search. Finally, alternative versions were to be possible without great expenditures of space.

Many text manipulations were to be possible, plus the ELF file operations (described in "A File Structure for The Complex, The Changing and The Indeterminate.")

Use of the DEC Display-8 was to make possible variable type-fonts and special characters in profusion, for all sorts of signalling purposes. This machine also makes text motions easier to implement.

Finally, XANADU was to meet the specifications of a general creativity console (see "creativity systems," another note), making possible spinoff and consideration of alternative versions and ready attachment of new hardware to fit reconfigured activity-spaces.

Appendix D
Transpointing Windows

To Ted Nelson Home Page
Back to "My Parallel Universe"

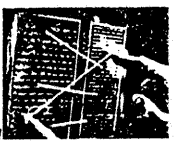

DOCUMENTS ARE PARALLEL BY NATURE

Documents are intrinsically parallel. Sound too complicated? Sorry. That's the nature of information-- it comes in packages which supplement and disagree.

Right now, most computer people still think of documents as independent and standing alone. This is a naive stage of understanding, looking at individual pieces and not seeing the whole-- like the biology of two hundred years ago.

But many documents are intrinsically parallel. This is where the hard document work lies, and what computers must be set up for.

To **EXAMPLES.**

Parallel documents come in different forms--

Parallel versions.

> My working definition of a document is this: a document is ***an arbitrary collection of versions having the same name*** (and possibly under control of the same owner), whose cross-connections and commonalities may be important. These are parallel.

Coupled documents.

> In addition, many documents travel together-- for instance, commentaries, which cannot even be considered without the documents which they comment on; or other documents implicitly connected, such as parodies and the original targets that they satirize. These too are parallel.

Non-parallel documents are the EXTREME CASE.

> The extreme case is a document with only one version-- with no variants, no alternatives, no history, no commentaries or other coupled documents. This is the trivially-simple case for which today's trivial text systems have been designed. And even though the large majority of documents may be like that-- alone and brotherless-- it is the more serious problem of multiple parallel versions that we must design for.

PHILOSOPHY OF HYPERTEXT

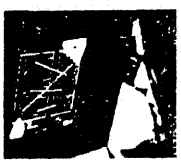

Maintaining, intercomparing and working with connections among parallel versions and coupled texts is what the hard document work is really about. This is what computer document work must be set up for.

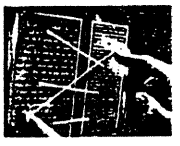

Back to "My Parallel Universe"

To Ted Nelson Home Page

Examples of Parallel Documents

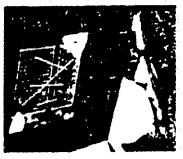

Back to "My Parallel Universe"
Back to "Documents are Parallel by Nature"

Documents are intrinsically parallel. Sound too complicated? Sorry. It's always this way. That's the nature of information-- it comes in packages which supplement and disagree.

Example: The Bible.

> The "book" that is called The Bible actually exists in many parallel versions. The Old Testament is basically the Hebrew Torah (of which an earlier, shorter version has recently been discovered in Ethiopia). The New Testament is a series of writings in Greek. Then there are the various translations of both Old and New Testament, either affected by culture and politics, or (some may say) guided by Divine Protection into perfect representations of the Word of God. Yet there are so many differing translations: for example, the English translations known as the King James Version (Episcopalian/Church of England), the Douay (Catholic) version, and various Twentieth Century versions. Then there are the translations into all other languages-- with differences in meaning that no one can know.
>
> All of these, taken together, are "The Bible."
>
> In addition, there are the Coptic Bible, the Greek Orthodox Bible, the Islamic Koran (said to be a parallel version), the Book of Mormon, and various books referred to as "Apocrypha", all of which have legitimate claim to be considered as versions of "The Bible". (A large work called *The Interpreter's Bible* collates some of these parallel materials.)

Would you like things to be simpler? So would a lot of people. Many people would prefer that you have access to only one version-- theirs. This is also the nature of information.

Example: "Rasho-Mon"

> The Japanese novel and Kurosawa film "Rasho-Mon" (-- note to my English-speaking readers: you have to pronounce it with the emphasis on the first O, or they won't know in Japan what you're talking about--) is the same story told six times by six different characters, including a ghost.

Which story is "true"? No one knows. But the the movie is about ***the parallel structure of all the stories together.***

Example: "Hamlet".

We think of Shakespeare's plays as being unique objects. However, Shakespeare's plays have come down to us in various versions, several of which are by Shakespeare, and many of which have been severely modified. "Hamlet" consists of *all* these versions, including also the "Hamlet" movies starring (respectively) Lawrence Olivier, Richard Burton, Mel Gibson and Kenneth Branagh.

Example: The Bill of Rights

When Americans refer to "the Bill of Rights", they are ordinarily referring to the first ten amendments to the U.S. Constitution. But there is another Bill of Rights from which it derived: the Bill of Rights of the state of Virginia, written by George Mason.

These two parallel Bills of Rights may be thought of, for many purposes, as one document.

Example of Commentary: Nabokov's *Pale Fire*

The novel *Pale Fire*, by Vladimir Nabokov, consists of two parts: a beautiful poem of a thousand lines, entitled "Pale Fire", supposedly by a character named John Shade; and a ridiculous commentary by a character named Charles Kinbote. The novel is the relation between these two parallel documents.

Example: VARIANT LISTINGS, e.g. Saints

Often two listings need to be compared. We may think of each of these listings as a version, and the set of listings as a document.

Example: The Catholic saints, the Episcopalian saints, the Byzantine saints, the Albigensian saints. These give us overlapping lists of the official saints recognized by each denomination. (And a time-series of Catholic saint listings, for example, would show the now-decanonized St. Christopher and St.Nicholas and when they lost official favor.)

Example: The Three Mile Island transcripts.

My favorite example of parallel documents comes from columnist/man-about-the-Valley Robert X.Cringeley, whose real name is Mark Stevens. He tells me this story from his work on the Presidential commission that investigated the Three Mile Island disaster.

According to Stevens, the Federal disaster management agency (whatever it was called at

that time) first denied having known about the Three Mile Island nuclear accident, and then revealed that ***actually they had had fifty telephone lines into the nuclear plant at the time of the accident– all fifty conversations being simultaneously recorded.*** Stevens says he had the recordings seized under subpoena and then transcribed.

We may consider these fifty recordings and their fifty textual transcriptions as ***ONE SINGLE DOCUMENT.***

Example: "The Carol Burnett Show" archive

It was recently announced* that 284 episodes of "The Carol Burnett Show", and all the associated documents (some 350,000 pages) will be donated to the UCLA library. We may think of this as being 384 parallel documents (a set for each show-- notes, scripts, and the recorded show itself); or we may think of ***the whole collection*** as a ***single, fairly large-scale parallel document.***

*(From http://cnn.com/SHOWBIZ/News/9811/18/showbuzz/index.html):

> LOS ANGELES (CNN) -- Comedienne Carol Burnett has announced that all 284 episodes of her "The Carol Burnett Show" will be housed for posterity at the University of California, Los Angeles. The collection, which also includes 350,000 pages of music, script notes and commercials, will be housed at the UCLA Library and in the university's Film and Television Archive.
> "I'm so happy that we can share the many years of music and sketches from our shows with the students of UCLA and people interested in the entertainment industry," Burnett said Monday in announcing her donation. "This is a wonderful way for the material to continue to have a life." The popular television comedy hour aired on CBS from 1967 to 1978 and won 25 Emmy Awards.

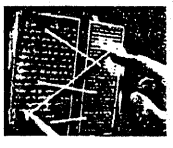

Back to "My Parallel Universe"
Back to "Documents are Parallel by Nature"

paraviz.html
98.11.29 (d9

To Ted Nelson Home Page
To "My Parallel Univrse"
To "All Documents Are Parallel"

PARALLEL VISUALIZATION: Transpointing Windows*

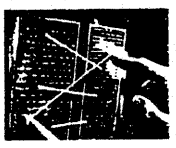

Here is the fundamental problem of hypertext: ***being able to see connections side by side.*** This is true whether you are looking at separate pages, different versions of the same document, commentaries on one document by another, or any other parallel document structure.

Here is the design I proposed in 1972 for a parallel visualizing editor. (This was long before the term "word processing" was invented, or the brain-dead function it denotes.) Current one-way hypertext, like the Web (and its predecessor HyperCard) lack this basic form of operation, which I consider to be the most fundamental tool of human thought. (Note that the photograph is a simulation: there were practically no office computers with screens in 1972; this is a cardboard model on top of a typewriter, with a celluloid picture to look like a screen. A nice friend of mine named Pat is the model.)

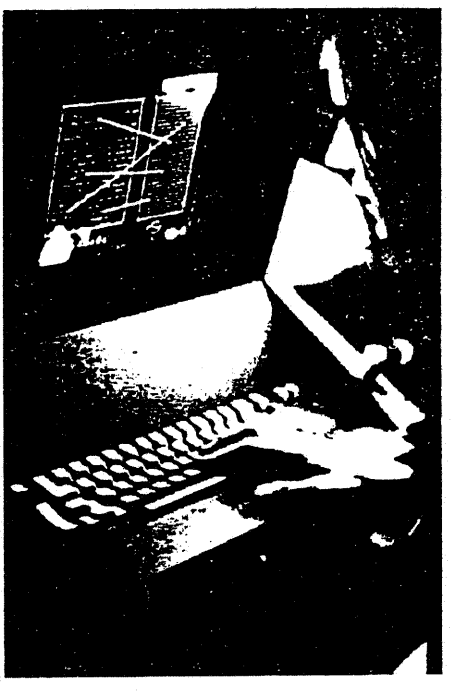

(Photograph from T.Nelson, "As We Will Think",
Proceedings of Online 72 Conference,
Brunel University, Uxbridge, England, 1972.)

This visualization is fundamental both for following connections and for intercomparing differences. Naturally, when scrolling either window, the endpoints of the connection-lines must continue to adhere to the elements

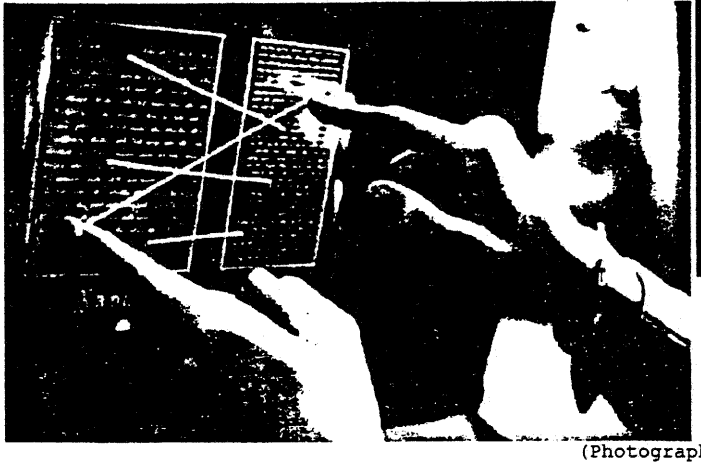

(Photograph from T.Nelson, "As We Will Think",
Proceedings of Online 72 Conference,

Brunel University, Uxbridge, England, 1972.)

they are pointing to, so the transpointing lines must move as required.

We now have a demo of this functionality, but unfortunately do not yet have permission from the vendor of the base software to distribute it to the public.

Note that this design was published in 1972, about two years before the unfortunate design of today's computer-screen windows was promoted out of Xerox PARC. (If I have any of these details wrong, I invite correction.)

---

* The term "transpointing windows" is from a much more recent publication: T. Nelson, "The heart of connection: hypermedia unified by transclusion", *Communications of the ACM* 38:8, (August 1995), 31-33. The term "Parallel Textface™" is still a trademark of Project Xanadu for products with this functionality.

To Ted Nelson Home Page
Back to "My Parallel Universe"

PARALLEL DATA STRUCTURES

Data should in general be parallel, with side-by-side indexing and markup. Side-by-side data structures make sense. (See reference 1.)

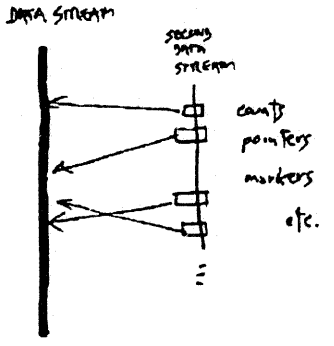

This makes it possible, for example, to use separate marking streams for different purposes, independently applying them to the same original data:

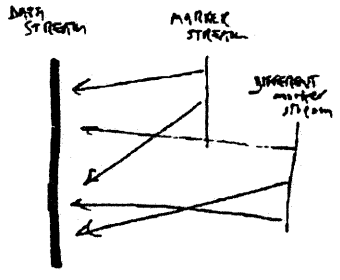

Unfortunately, the tradition is to jam the data all together, and the result is our messy file systems-- and worse, SGML and HTML. The currently popular model of SGML (and its derivatives HTML, XML) scrambles together two intrinsically different kinds of data which are intrinsically parallel. This mixing makes it impossible to switch between markup structures (as in the second illustration, above.)

Examples of parallel data in popular software

Some people have seen the advantages of parallel data structures. One nice example is in the popular Eudora e-mail reader. The PC version of this program maintains the original Unix file structure of concatenated emails, but then marks their beginnings (and other information) in a

parallel stream, the .toc file (table of contents). This is a nice example where the addition of a second parallel data stream allowed the first to be untouched.

A better-known, but less successful, example is the distinction in Macintosh files between "data fork" and "resource fork". Unfortunately this distinction has ceased to have any coherent fixed meaning, and whatever it was originally intended to clarify has long since been muddied.

Deeper uses: The Xanadu® software family

The Xanadu family of data structures has always (at least since 1970) been designed around parallelism between contents and markup streams. (We expect to publish the original design papers hereabouts shortly.)

Reference 1. T.Nelson, "Embedded Markup Considered Harmful", In *XML: Principles, Tools, and Techniques* (*World Wide Web Journal* 2:4, fall 1997). This paper is also on the Web at http://xml.com/xml/pub/w3j/s3.nelson.html.

Appendix E
Xanadu User's Manual, 1976

USER'S MANUAL
for the
XANADU INFORMATION SYSTEM,*

XANADU HYPERTEXT NETWORK*
and
XANADU INFORMATION SERVICES.*

* Note: "Xanadu" is a trademark and service mark for computer products and services offered by Theodor H. Nelson.

XUM-1 Version 1, 23 November 1976.

©1976 Theodor H. Nelson

Dedicated to the gallant freedom fighters of Project Xanadu, especially William Barus.

PRICE $5.

The Xanadu Information System is a computer-based information system of the hypertext type. That is, it allows reading and writing essentially like the reading and writing we have always done, but with nearly instantaneous access to writings, and mmwith multiple attachments and linkages possible between the writings.

This manual describes operation of the system as planned for August, 1977 in an experimental version. This version of the Xanadu Information System, and its component Xanadu Hypertext Network and Xanadu Information Services, will be designated XX-1 (Xanadu, experimental, version 1).

All references to the Xanadu Information System refer to the XX-1 version, and these specifications are subject to change without notice.

ALL DATA IS STORED AT YOUR OWN RISK.

The Xanadu Information System consists of one or more special computer terminals, which may be anywhere, and one or more computers, with storage, at one or more central sites. The computers and interconnections constitute the Xanadu Hypertext Network, and their operations are the Xanadu Information Services.

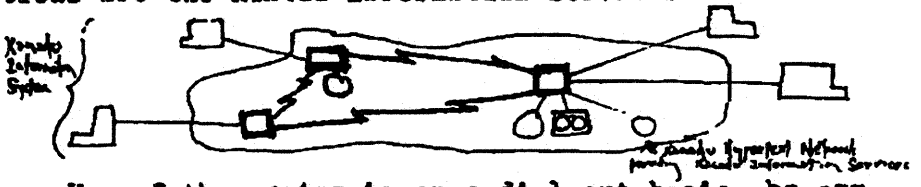

Use of the system is on a dial-out basis, by appointment. Communication between terminals and the network is in full-duplex ASCII code (see Appendix A, "Communication Protocols").

Users pay for connection time, but not for ordinary processor services or mass storage access. Included in the base rate is cumulative perpetual mass tape storage. The allowed accumulation is set to a moderate typing speed. The base rate also provides for the mounting of one tape per hour. Also included are contractual royalty payments to owners of copyrighted materials which are publicly available on the system.

Additional storage, additional tape mounting, increased accumulation speeds and the carrying of material on disk, all are available at additional rates (see "Tariffs," Appendix B).

Other vendors are encouraged to supply similar compatible services, and under compliance with certain specifications may be licensed to use certain trademarks (see Appendix C, "Trademarks and Licenses").

THE TERMINAL

The terminal initially employed in the Xanadu Information System is a small computer with a color screen and keyboard, offering both extensive text display and some graphic capability. We will refer to it here as the color terminal, or as Xanadu Terminal CT-1™ (See Appendix D, "Recommended Equipment.")

The color terminal consists of an Intecolor® terminal computer (manufactured by Intelligent Systems Corporation of Chamblee, Georgia), with various specific options: 48-line text display, lower case, graphic capability, 16K RAM, and our own Xanadu Terminal Program™ in a special read-only memory. It is used with ordinary modem and coupler to connect to the Xanadu Hypertext Network over ordinary telephone lines.

The keyboard is a standard typewriter keyboard with (in addition) numeric pad and certain color keys, which are used in the operation of the system.

SYSTEM CONTENTS

The system holds "documents," which may be either ordinary written documents, or certain new forms permitted by the system.

A document may contain written text, including headings, tables of contents and indexes. No provision exists for illustrations at the present time.

It may also contain <u>links</u> to parts of itself, or to other documents.

> It may contain jump-links to itself or other documents. (A jump-link, symbolized here as ▣, means that another piece of writing is connected, like a conventional footnote. The jump-link allows the user to jump immediately to the connected matter.)
>
> It may have multiple links to another document ("collateration"). Collateration, which we symbolize here with ▤, is roughly the equivalent of marginal notation or side-by- side printing. Collateration allows the reader to view corresponding or related parts between documents.
>
> It may embrace other documents, or parts of them, permitting instant viewing of the embraced document as part of the present one. This we call "quote-windowing," symbolized by ▣. A quote-window, in which part of another document is viewed, may be expanded instantaneously to a view of the quoted material in its original context.

Finally, a "document" may consist of no text, but only links of the above types among previously stored texts. For instance, an "anthology" may consist simply of quote-windows to many different texts stored within the system.

4

PRIVATE VS. PUBLIC

A document may be private, meaning that its use is restricted to its owner and those to whom he grants ppermission.

A document may be public, meaning that any user of the system may have access to it.

Not monly that, but any user of the system may, of his own choosing, create links into any public document, embrace it within other documents, and so on.

Rendering a document public is not reversible. It is the intent of the system that everything which has been made public remain so, regardless of what change of heart the owner may have. Once the public has access to a document, it has permanent access to that document.

COPYRIGHT

(A new copyright law has just been passed by Congress, and will go into effect at the beginning of 1977. The following remarks may require revision in the light of that new law.)

Copyright is a part of the Xanadu System. It is intended that the Xanadu Information System will respect and preserve the copyright of all owned documents, whether public or private. However, due to the notorious fallibility of digital computer methods, no guarantee can be made that this will function correctly, and the inadvertent loss of copyright may be thought of as a danger of the system.

Copyrighted materials will be displayed with a visible copyright notice at the beginning of, or in association with, each presentation.

A royalty fee will be held in trust for any author of an owned public work whose work is viewed by someone else, excluding persons involved in system maintenance and administration. Royalty arrangements for private documents are also possible.

The maximum royalty of 10¢ per hour (an experimental rate for XX-1) will be paid to the owner of a work which is viewed exclusively and alone. If more than one work is on view at a given instant, that fee will be divided among the works in proportion to the amounts shown. All royalties will be prorated to the second of time. If a pproportion of what is on the screen is not copyrighted, or is owned by the viewer, the owner of copyrighted material will nevertheless receive royalties in proportion to all materials shown. For instance, if three-tenths of the material seen in an hour is copyrighted, five cents will be paid in royalty. These proportions will hold regardless of the amounts viewed; skimming and fast reading will cost the same as slow usage.

Joint authorship is of course permitted. When joint authorship is specified, royalties are held in an undivided account for disposition according to the owners' intention.

Anyone may waive royalties, or designate another recipient.

It will be noted that these arrangements are part of a private contract between the provider of these services and all parties using the system. Disputes as to royalties will be decided by the administering party until these matters can be vested in an appropriate representative body. (See Appendix E, "User Contract.")

OPERATION

We will first describe typical operation of the system; then certain xbackground views and operations; then the way that documents are looked up and created.

TYPICAL OPERATION

In typical operation the user has two texts on the screen. Each of these may be in more than one color combination, as chosen by the user. For each of several related documents the user may choose a different color, or color combination (such as white lettering with no background, red lettering with a blue background, black lettering with a red background--some are more readable than others). A residual color combination (such as white on black) is used for any other texts.

Any two collaterated documents may be presented on the screen side by side, each in the color the user has chosen for it.

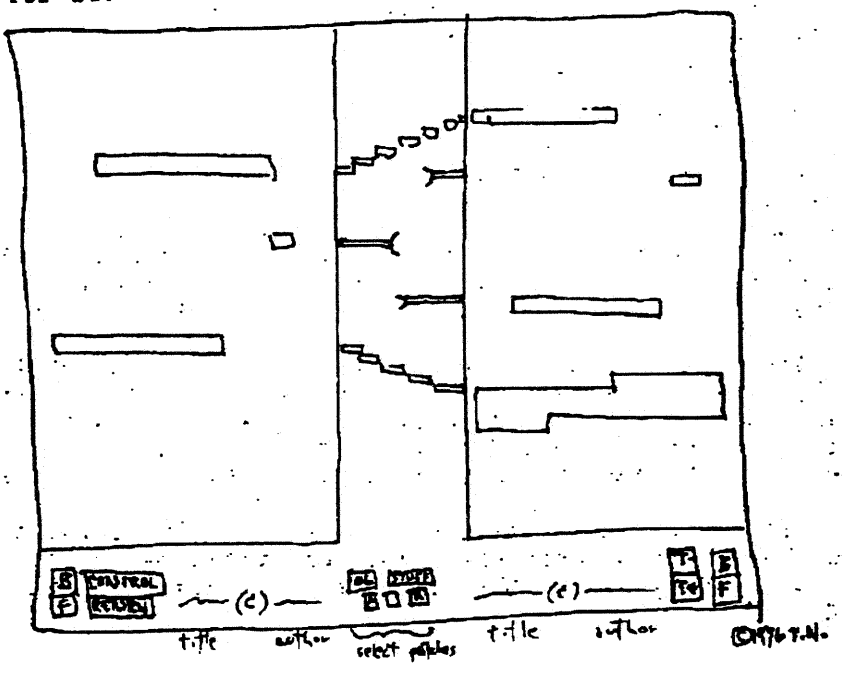

6

The independent text, which may be moved freely (see Computer Lib, p. 76) is chosen by the user implicitly when he tells that text to move (described later).

If a collateral link exists between a portion of the independent text which is shown on the screen and a portion of the dependent text which is shown on the screen, a line is seen to connect them.

Note that since the color terminal does not permit graphics in the same spaces as text, the collateral link-lines are shown relative to the ends of the lines of text. The portion of the independent text which attaches to the collateral link appears field-reversed, with the background color in the foreground and vice versa.

If the link is from one text to an unshown part of the other text, it appears as terminating in a parenthesis. The user may point at this line (with the cursor) and press HOME to make the linked text appear.

If other documents are simultaneously collaterated to either text, their collateration links are shown in the colors specified for those other documents.

If the link is from a shown text to another text, the user may point with the cursor and press HOME to make the linked text appear.

(The user may hold non-contiguous sections of dependent texts on the screen together, from one or more documents, by typing an ampersand between one collateral jump and the next.)

Available footnotes or jumps to other documents are indicated by asterisks/of the other document's color. or link-lines

To take such a jump, the user moves the cursor to the desired asterisk or link-line with the arrow keys, then hits the HOME key. The linked material then appers.

The material jump ed to appears on the side of the dependent text, obliterating part or all of it. If the material is a short jump-linked piece, such as a footnote, it overlays a part of the dependent text, with a black area above and/or below, and a line of asterisks running from the jumpout-point in the independent text to the beginning of the material. The asterisks are in whatever color has been chosen for that text, or in the residual color.

The bottom line is normally used for control functions, simply because the text can be scrolled up and the control patches unobtrusively and conveniently regenerated in that position. Some control functions spill to the line above, others are "hidden" in an implicit line "below" the screen.

Brief title and copyright notice (with initials) are shown beneath each document. The full title or author may be found by moving the cursor to that area and hitting HOME.

Either text may be moved by moving the cursor to the F (for Forward) under that text, and hitting HOME. The B moves it Backward. (Rates may be adjusted.)

To take a jump, the user moves the cursor to the selected asterisk or incomplete link, then hits HOME.

To return from a jump, the user selects RETURN and hits HOME.

Colored patches in the middle of the control line (with an initial or brief description) signify other current texts. To go to another text, the user moves the cursor to one of these patches and presses HOME. An arrow key to left or right selects the desired side of the screen. (Ordinarily this selects whatever part of the new text corresponds by collateration to the text which remains showing. (If no determination can be made, the jump is to the beginning.) If this method is used to move a text from one side of the screen to the other, the same position is maintained.)

When the user wishes to make an editorial change, he being the owner, he drives the cursor to the bottom line and hits the down-arrow, causing everything to scroll upward, including the bottom control line; an edit-line appears below:

[* LINK] [= LINK] [" WINDOW] [COPY] [MOVE] [DELETE]

The leftmost three patches allow the creation of links, or, if preceded by MOVE or DELETE, the moving or deleting of links. COPY, MOVE and DELETE may refer to text as well. Up to three colored file-select patches also appear, to prevent ambiguity as tom which documents are referred to when link-changes are made.

Insertion is implicit: merely typing in creates an insertion, if the edit-line is showing.

The insertion of an exclamation point, preceded and followed by a space, creates a cut-point for the MOVE instruction (as in the JOT system; see Computer Lib, p. 79). Three or four cuts determine a rearrangement.

Hitting the CONTROL patch causes a jump to the CONTROL face.

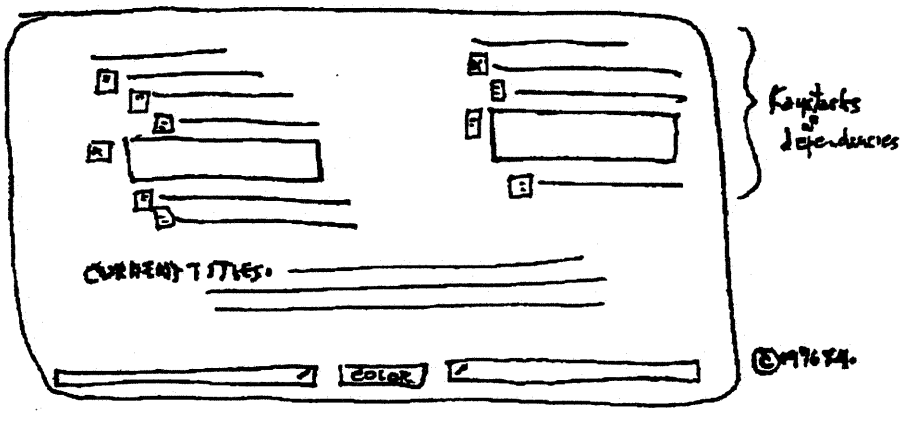

This shows the current dependency-status of the two texts on the screen (in Kay-stacks); the current titles in their selected colors; and two width selectors.

The dependency-chains allow the user to see, instantly, the status of each document he is currently viewing, and how he got there. The topmost title indicates (in its color) where he started; each succeeding line shows what he jumped to next, and its relation to the previous. The current document is shown blinking.

Below the current document may be seen titles of documents which have been x jumped to and returned from, after the style used by Alan Kay for his SMALLTALK windows ("Kaystacks" ; see Kay et al., "Personal Dynamic Media.") Note that titles under the blinking one do not reliably show dependency or order of recency, but only size of variable-depth stacks which may h have been visited independently.

To widen or narrow a viewing panel, the user may seize the slash at the end of a width-stripe at the bottom, hit HOME, move it sideways with the arrow keys, and clinch the new position with another HOME.

Hitting the COLOR patch allows the user to designate new colors for his chosen documents.

Driving the cursor below the bottom line causes the display to scroll and brings the LOCKUP LINE, allowing selection from all documents in the network.

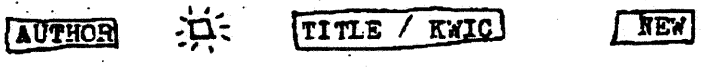

The cursor appears between the words TITLE and AUTHOR. Then user designates which by an arrow key; then types as many letters of the title or author's name as he chooses, followed by HOME. (The user may also seize and position an arrow for an approximate jump into the list.) If the user selects KWIC, and types a specific word followed by HOME, the titles containing that word will appear, arranged with that word in a column.

Public files are listed black-on-white. The user's private files are listed blue-on-white. Other private files, open to the user, will be listed red-on-white.

The cursor is used to select the file (document) desired, followed by HOME. The user now selects a color-pair for the present showing of that file. Driving the cursor against the bottom line serves to scroll through the catalog.

This continues as long as desired.

The user may now return to the Control Face, by sending the cursor through the top line.

To create a file, the user selects NEW and supplies a title, selecting colors at the same time. (These colors are temporary and will not be permanently associated with the document.)

Appendices

9

A. COMMUNICATION PROTOCOLS

B. TARIFFS

C. TRADEMARKS AND LICENSES

D. RECOMMENDED EQUIPMENT

E. USER CONTRACT

Appendix A. COMMUNICATION PROTOCOLS

While a final determination has not been made, a system of encoding the functions described here will be furnished, permitting users to adapt their own programmable terminals for operation with the network, and permitting other suppliers to offer similar or compatible services under our trademarks and licenses.

These codes will be selected from the American Standard Code for Information Interchange (ASCII), used in full-duplex mode.

Functions specifiable will include:

Forward
Backward
Jump #
Jump Collateral #
Open to Window Context #
RETURN
Insert
Delete
Cut
Move
Backward in time (one step)
Forward in time (one step)
Backward in time (one version)
Forward in time (one version)
Tree select (in revision tree)

10

Appendix B.

XANADU TARIFFS, System XX-1; estimated, November 1976.
Note that these prices will be subject to day-to-day changes without notice, as well as changes by time of day.

It will be seen that not only are these charges very different in rate-structure from conventional computer time-sharingx systems, but they are subject to absolute control by the user, based only on time of use and services explicitly requested.

It will also be seen that the fast reader or skimmer pays the same amount, per time, as the slow.

Of course, speed of performance is not specified here. How fast the system works determines the effective prices of the service, since our service prices are fixed per unit time. If services are slow they may be said to be, say, twice as expensive as they would be if they were twice as fast.

The intent is of course to drive delays down to the practicable minimum.

BASE CHARGE:

Connection, full usage and access to all public on-line materials; unlimited jumps; unlimited editing; unlimited placemarkers and rovers and faces, subject to: tape storage equivalent of 5 keystrokes/second over time of use; mounting of one tape per hour or fraction thereof; including royalties to authors up to a total of 10¢ hour.

$2/hour

Communicat ion cost	Standard phone cc. rates
Additional storage	10¢/thousand characters
Mounting of tape cassette (additional to one per hour allowed in base rate) (Note: this is subject to availability of tape drives)	10¢/mount
Standing disk storage	$1/thousand characters/month
Mountable disk storage, floppy (subject to availability of drive)	10¢/mount
Dedication of floppy	$1/hour
Temp orary disk load, by prearrangement	10¢/thousand characters/hour
Preludal mount-and-load (anticipatorily moving key-structures to disk)	$1

tape redundancy	50¢/hour
triple tape redundancy	$1/hour
PRINTOUT (Diablo)	25¢/page
Note that this is subject to a contractual royalty to all copyright holders.	plus royalties
Base price includes no carbon copy, 1-color carbon ribbon, no overstrikes or graphics.	
All other printout	estimate on request

ADDITIONAL SERVICES

A number of additional services are planned, for some of which there will be additional charges. These include historical backtrack, viewing histories, document scans, document concordances, running-concordance edit mode, concordance KWIC viewing, and so on.

Appendix C. TRADEMARKS AND LICENSES

After the period of experimental (XX-1) operation, other suppliers will be encouraged to offer compatible services, and a number of trade and service marks will be made available, under licenses, to those supplying compatible or official Xanadu services.

These will include:

XANADU™ XU™	Xanadu Products and Services
XFS℠	Xanadu Full Service
XM℠	Xanadu Mobile Service
XCA℠	Xanadu-Compatible Archive
XA℠	Xanadu Approved Accessories
XT™	Xanadu Terminal
Porta-Xat™	Portable Xanadu Terminal
XNE℠	Xanadu Network Entry (for time-sharing systems)
XA℠	Xanadu Access
XED℠	Xanadu Editing Services
XUL™	Xanadu Ultimate Language
XC™	Xanadu Center
Xanfolk℠	Xanadu Employees

These will be used in conjunction with certain styles of "X" and "Xanadu," e.g. the X-Hand:

The principal conditions for such licensing will be our examination and verification of goods or services; and conformance of these goods and services with the general copyright and royalty arrangements native to the Xanadu Information System. License fees will be low.

Appendix D. RECOMMENDED EQUIPMENT

Note: in recommending this equipment, we make no guarantee that the services described herein can be feasibly offered on a useful basis. Any decision to purchase such equipment must be based on the purchaser's own estimate of the usefulness of this equipment in his or her general computer pursuits, of which the services described here may or may not become a feasible part.

Xanadu Terminal CT-1™: Intecolor® terminal with options described earlier, available from the manufacturer or Itty Bitty Machine Co., Evanston, Illinois. Price in the vicinity of $4000.

While this unit is also available in kit form, we do not recommend it.

Appendix E. USER CONTRACT

Users of the system will be subject, implicitly or explicitly, to a service contract embracing liabilities of the supplier and their limitations, and copyright arrangements along the lines described here.

It is intended that detailed guidelines for the royalty system will grow out of usage, and that an independent body representing service organizations and user constituencies will be created to adjudicate copyright and royalty disputes within the system.

Appendix F
Walkthrough of the Xanalogical Model

ApF-xu88walkthru
02.01.10 (d26

To Ted Nelson Home Page

Appendix F.
Walkthrough of the Xanalogical Model.

This is excerpted from "**Xanalogical Structure, Needed Now More than Ever: Parallel Documents, Deep Links to Content, Deep Versioning and Deep Re-Use**"
at http://www.sfc.keio.ac.jp/~ted/XUsurvey/xuDation.html

Theodor Holm Nelson, Project Xanadu* and Keio University

THE XANALOGICAL MODEL: MECHANICS AND FORMER SECRETS

Thus far we have considered two levels of this proposed literary system: intended screen appearances, and the two complementary forms of xanalogical connection (content links and transclusion). Now we turn to its supporting system of indirect representation by version lists and connection lists.

What will be presented here is a fundamentally simple model for a large-scale hypertext literature built on xanalogical structure, extremely different from the World Wide Web, intended to support profuse fine-grained 2-way unbreakable links and principled re-use.

This underlying model is simple but generally not known, in part because our methods were under complex proprietary ownership, and thus trade secret, until the open source release of prior Xanadu code in August 1999. However, these proprietary methods (especially enfilade theory and enfilade specifics (26)) were really for efficiency in carrying out methods like those to be described below.

The central proprietary secret this all relied on-- which we considered whimsically obvious but never stated publicly-- was the freezing of content addresses into permanent universal IDs, below.

We will go over our fundamental model, resolved to specifics of the xu88 design (simplified for explanation). By being this specific we can hope to be tediously clear. While xanalogical structure is conceptually independent of this more detailed model, the xu88 design serves as an existence proof and an available client-server ensemble, now available at Udanax.com under an open source license with the names Udanax Green and PYXI, and thus constitutes (at the risk of some confusion) what is usually called a "reference version" of these concepts.

USE IN PLACE: REFERENTIAL FLUID MEDIA

In this model, documents, versions and links are structures which point to content and use that content by reference. Operations on these structures of lists do almost everything. This is the opposite of the prevailing, and extremely limited, embedded approach (**108**). What follows will be a detailed walkthrough of data and methods for implementing xanalogical structure with this referential data model.

Other people's text systems typically bring text into a software buffer and modify it-- thereby losing all information as to origins, other versions or other connections.

Our alternative is referential editing of referential documents. We give any text (including newly-input text strings) a permanent address, and thereafter use that text by pointing to it. Thus text is used "in place", accessible simultaneously to many possible applicative uses. The same method is used for other data consisting of countable, separable media elements. (Such reusable media content with permanent element addresses we now call "fluid media".)

Let us see how referential operations work for bodies of sequential text (Fig. 7), then consider links and transclusions. Discussion of richer structures is for elsewhere.

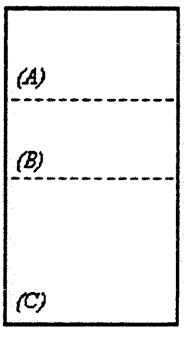

Fig. 7. Sequential text as it often appears on a screen. (Parenthetical labels and portion boundaries are only for matching to fig.8.)

A sequential document or version is represented by a content list, which is the fundamental form of representing or transmitting a document. (Note that the operative unit of the xu88 design is the *version*, but the term "document" will be clearer to most readers.)

Though the content may appear on the screen as an ordinary block of text, structurally its chief representation is a list of contents-- a sequence of reference pointers. These pointers designate spans of characters (or other elements) (fig 8). These spans, concatenated to make a virtual stream of the designated elements, thus represents the sequential text of fig. 7.

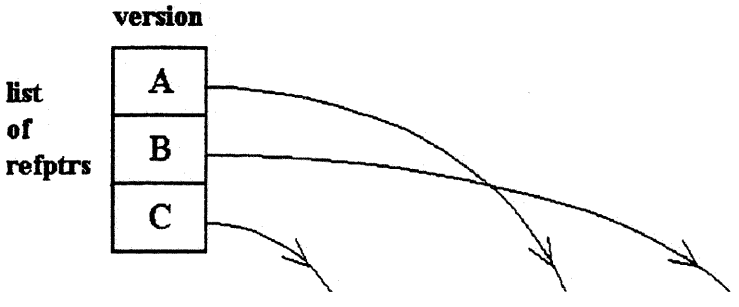

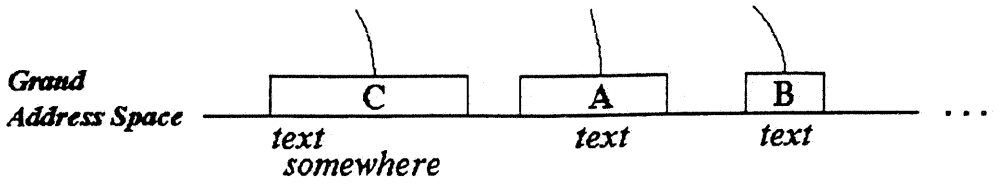

Fig. 8. List of contents representing the text version of Fig. 1. Proportions are not significant.

All editing operations change these content lists, which remain the fundamental form of representing or transmitting a document; and all search and editing operations are on these content lists. These operations are identically performed, regardless of who owns which parts of the content or where the content actually resides. Pulling up the actual contents for user viewing or editing is essentially cosmetic and informative to the user, since internal operations are on the version lists themselves.

These spans of content may be in separate locations anywhere in the address space (i.e., the universe). But content may be cached in many places, as long as it has the same permanent IDs. (However, this is a level of indirection separate from the basic model--a provision which onlookers have generally overlooked, to their great confusion).

Hollywood discovered this method separately. This is how movies are now edited; this is how the xanalogical model has always edited contents of any countable kind. What Hollywood calls an EDL, or Edit Decision List, we call a content list or version list.

Thus the delivery of a document is in two logical phases: the content list, then the fulfillment of the contents. Getting the actual contents is like going to retrieve the film negative itself from the vault-- conceptually a final stage. Having two stages may sound like unnecessary overhead, but it is similar in some ways to the many stages of protocol exchange in the delivery of a multipart Web page today. It may also be compared to the overhead of the Internet itself-- enormous in terms of program cycles and data packets, but an enormous simplifying benefit, rather than a burden, in its result.

CONTENT LINKS, AND HOW THEY SURVIVE

The xanalogical content link is not embedded. It is *applicative*-- applying from outside to content which is already in place with stable addresses. Xanalogical links are effectively overlays superimposed on contents. Any number of links, comments, etc., created by anyone anywhere, may be applied to a given body of content. By this method it is possible to have thousands of overlapping links on the same body of content, created without coordination by many users around the world. (Today's embedded links, as in HTML-- a later and cruder design-- cannot be profuse, overlapping, bidirectional, or diversely created in this way (**108**).

Each xanalogical link is implemented as a separate unit which connects spans of content. This in turn means that a link can apply to any of that content, wherever the content may be re-used.

First let's consider a simple document, then a comment link on it, and watch how that content

link continues to adhere to the changing document. (We will assign toy addresses, simpler than those in the xu88 protocol, to make the concepts clear.)

Suppose author Adam creates document A, version 1, which he publishes on line as the first document in this format. He publishes all 300 characters at once, causing these characters to be given consecutive universal addresses. The text contents of document A v.1 are therefore registered as having consecutive addresses from 1 to 300.

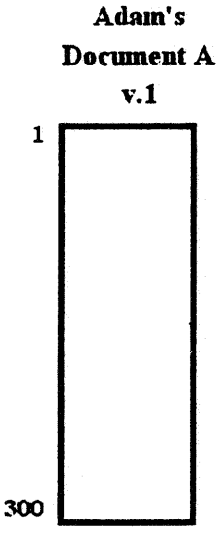

Fig. 9. Screen visualization of Adam's document A, still an unbroken stretch of content bytes. (Numbers on the side are for following further examples.)

This document is internally represented as a table of pieces. Since it so far only has one piece, it is saved as a single reference pointer to the consecutive characters 1 to 300.

Adam's Document A

v.1

list of ref. ptrs (only one)		**start**	**end**
	ptr 1	**1**	**300**

Fig. 10. Internal representation of Adam's document: a list consisting of one reference pointer, since it references a consecutive span of content.

THE CONTENT LINK

Now along comes a second author Barker, who publishes document B, version 1. The text contents of document B v.1 are registered as 217 consecutive bytes with permanent addresses 301 to 517, also an unbroken stretch of content bytes.

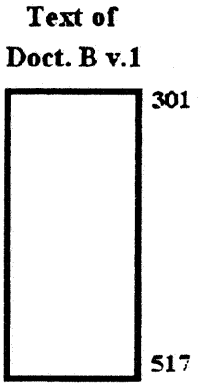

Fig. 11. Screen visualization of text portion of document B v.1. (Addresses shown only for clarification.)

Now for a link.

Suppose that in the text of Barker's document B, he has made a comment on Adam's document A. The text of the comment consists of characters 351 to 400; they comment on characters 101 to 200 in the middle part of A. While Barker's comment already resides in his text in written form, it gives no direct access to the material being commented on. So let us say that as the author, Barker chooses to make this textual relationship explicit, representing it as a comment link between those two spans (presumably using some comfortable visual editor to select the spans and specify the link type).

The link is from an entire span on the first side to an entire span on the second side. (Note that anyone from anywhere can make a content link to any published content.)

One way this might look on the screen: a stripe between both documents (fig. 12).

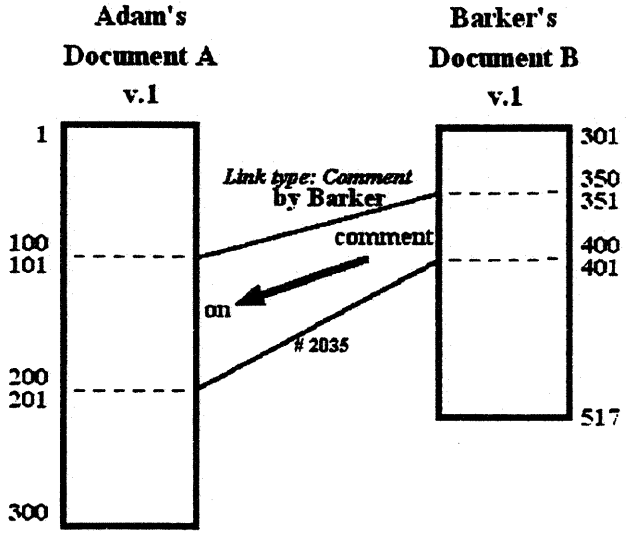

Fig. 12. Possible screen visualization in transpointing windows of

comment link (Barker's link, element 2053) connecting text from document B v.1 to document A v.1. (Text addresses and dividers shown only for clarification.)

Content links are first-class, free-standing, and addressable. Even though the comment is a part of Barker's document B, the link representing it is a distinct entity. Let us say it is somehow assigned address 2053 to stress its utter independence of the text of document B.

The internal structure of the content link has three parts (fig. 13): a list specifying the first side, a list specifying the second side, and a type designator. (These are called in xu88 the 1-set, the 2-set, and the 3-set.)

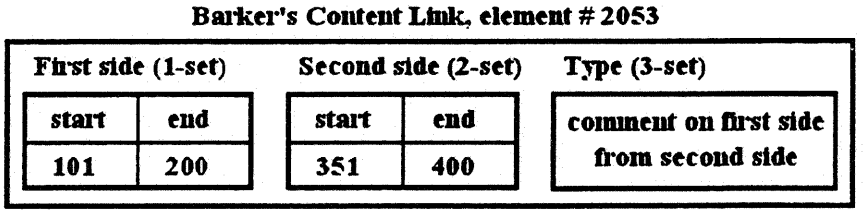

Fig. 13. Internal representation of Barker's content link 2053 from document B v.1 to document A v.1. The content link connects all elements on the first side to all elements on the second side.

Making the link is one logical data step, including it in a document (fig. 14) is another (although the visual editor should usually combine these steps from the author's point of view).

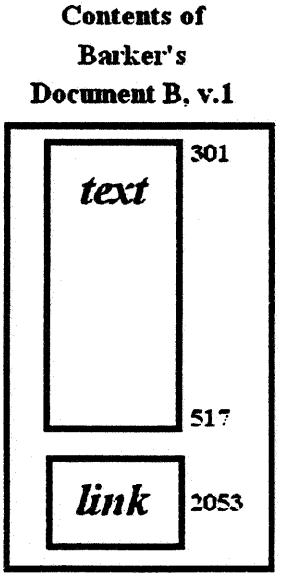

Fig. 14. Virtual contents of Barker's document are the text and the link.

The link is put into the document as another pointer in the document list of B (fig. 15). Therefore the document is represented as a list pointing to B's characters *and to the link*,

which is an enumerable element like the characters themselves.

Barker's Document B

v.1

list of ref. ptrs

	start	end
ptr 1 *(text)*	301	517
ptr 2 *(link)*	2053	2053

Fig. 15. Content list of Barker's document B, version 1.

Although this link is a part of Barker's document (fig.14), it is also externally addressable. (Thus the link may also be made part of any other document by transclusion-- like any other content, as will be seen.)

HOW LINKS DON'T BREAK

Now the fun begins.

Suppose Adam now publishes version 2 of document A. To make version 2, Adam inserts 67 characters, which are assigned addresses 7784 to 7850. This breaks the original reference pointer in two, and splits the content list of document A into three reference pointers.

Adam's Document A

v.2

list of ref. ptrs

	start	end
ptr 1	1	150
ptr 2	7784	7850
ptr 3	151	300

Fig. 16. Internal representation of Adam's document A version 2: the first pointer has been split into two, with an insertion between.

Now what happens to Barker's link? The link itself has not changed, but the pattern of its connective positions on the new version has changed, so it will look different. In document A v.2, although an insertion has made into the middle of the previous text, all the characters from document A v.1 which participated in the link-- that is, were commented on-- remain. Thus the system can show the link to all these characters in their new positions. One possible visualization is seen in fig. 17.

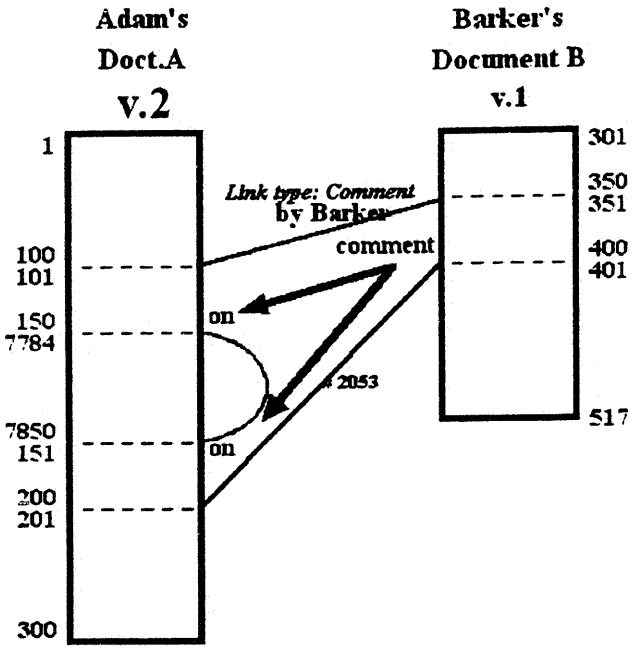

Fig. 17. Screen visualization in transpointing windows of link 2053 still adhering to the same characters in document A as rearranged in Adam's version 2. (Addresses and dividers shown only for clarification.)

Now again Adam edits Document A, creating a shortened version 3 by deleting characters 101 to 166. The resulting document content list is now

Adam's Document A

v.3

list of ref. ptrs

	start	end
ptr 1	1	100
ptr 2	7784	7850
ptr 3	167	300

Fig. 18. Internal representation of Adam's document A version 3: content has been deleted from version 2.

It happens that some of the characters that Barker originally commented on have been deleted. In this new version, fewer elements in the destination span still exist-- only characters 167 to 200. But link 2053 remains attached to all of these surviving characters in any new version. One possible visualization is in Fig. 20.

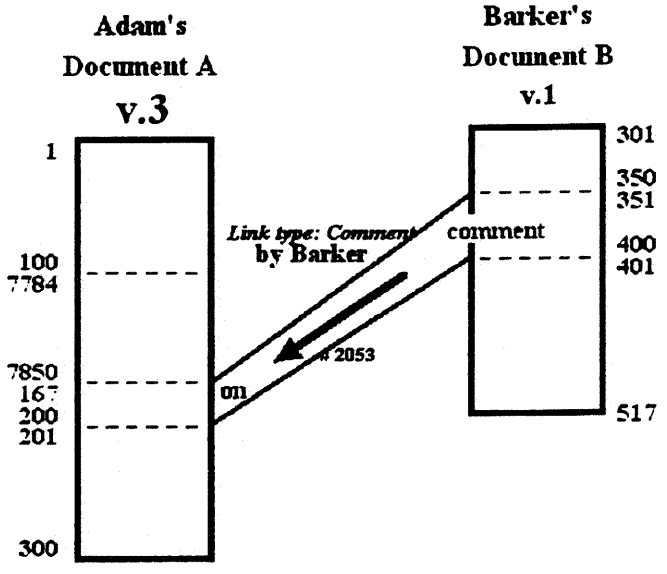

Fig. 19. Screen visualization in transpointing windows of link 2053 still adhering to remaining linked characters of Adam's version 3. (Addresses and dividers shown only for clarification.)

MECHANICS OF LINK DISPLAY

Let us clarify the mechanics of how the link is shown.

To make the screen presentation of link 2053 between two versions, the system compares content link 2053 against the content lists of whatever specific versions are being shown. The system finds those characters in each document on the screen which participate on either side of the link, and displays them on the screen with connection beams.

The generalized method of showing a content link, given two elements or versions chosen to be seen in transpointing windows is to find all the addresses of elements which are shared between the content link and the specific versions, by finding all the overlapping addresses of elements which are shared by the spans of the versions and the link; and show connection beams as appropriate. (If no shared elements are present on one side, and the user wishes, the user may also go back to some previous version that still had shared elements.)

POWERS OF THE CONTENT LINK

It will be evident that this mechanism-- keeping a link attached between occurrences of surviving content through successive versions-- will retain structure that otherwise will be lost. Just as we cannot expect Adams to coordinate with Barker, we cannot expect any rich coordination among the ever-changing electronic documents of the world. So we represent the literary connections that move, or remain.

Obviously there can be different problems with this, depending on the particular revisions of

different documents, documents' semantics. Still, this method manages the *presumptive* inheritance of meaning among versions; and the reader's intelligence, plus access to other versions, should take care of the rest.

Example: detailed marginal notes. These tend to lose their attached meaning as soon as new editions of original document come out; trying to maintain detailed marginal notes on the Web is virtually hopeless. But use of the method for annotation, in future formats, will give some assurance that the annotations will have a much longer life expectancy.

This approach is extremely different from the 1-way links of HTML, each of which can have at most one successful destination (with tangled local and cached exceptions). The HTML link is attached not to a document but to a virtual place-- a file (or marker in a file) in a specific directory on a given server, not a document. The target is deemed to be whatever is at that address, if anything. It is like a shop-window-- we see it however it may be currently decorated, or boarded up.

(How to show many overlapping content links is another issue: color, translucency, reduction to lines are among the possibilities; but these are "interface issues".)

GLOBAL LINK FOLLOWING

But how, then, can a content link be followed into the blue, like an HTML link with unknown destination, when it can be satisfied in principle by any document or version having those specific content elements? The answer is for the user to pre-reduce the search-- making some contextual selection, such as what author's work or what collection to look at, whether to ask for the latest version or the original. Then, when the system says there are many possible hits (like a Web search engine), refine again (as with a Web search engine). This may present an interface challenge, but so does every rich capability.

Users are already accustomed to sorting for relevance among millions of Web pages; the prospect of sorting for relevance among millions of links should be no more difficult conceptually. We believe that *recommendations of links*, a new way of effectively creating moderated newsgroups of content, will become an important genre.

FINDING TRANSCLUSIONS

Transclusions, or visible re-uses across literary boundaries, may be found straightforwardly with the referential model. We find them by address comparison.

Consider Adam's three versions of document A. What content do these versions have in common? We may find out easily by intercomparing the version lists of document A versions 1-3 (figs. 10, 16, 18), and displaying them with transclusion beams (fig. 20).

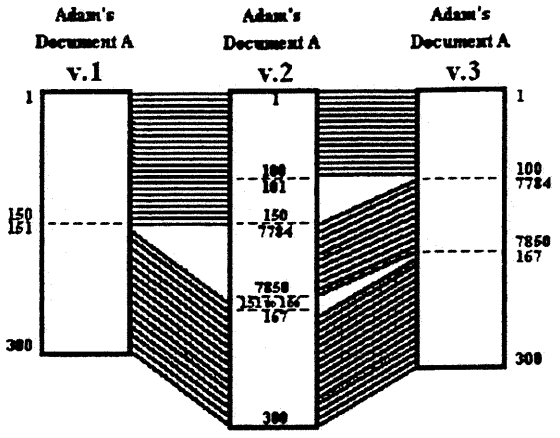

Fig. 20. Transclusion between versions as it might be seen in transpointing windows, showing identity of content across literary boundaries. Note that fluid-media transclusions are always the same length.

Any transclusion may also in principle be made "clickable", so that a user viewing a quoted part of some document may pull up the original context in a transpointing window. The system does this by finding the quotation's position in the original version, opening the original version to that position, and displaying it with transclusion beams.

If this seems too abstract, consider a familiar example. In emails, we often quote previous emails-- often marked by angle brackets on the quoted portions-- but the context of each quotation is lost. A transclusive email viewer (assuming supporting mechanisms) would maintain the connection to the original context (fig. 21).

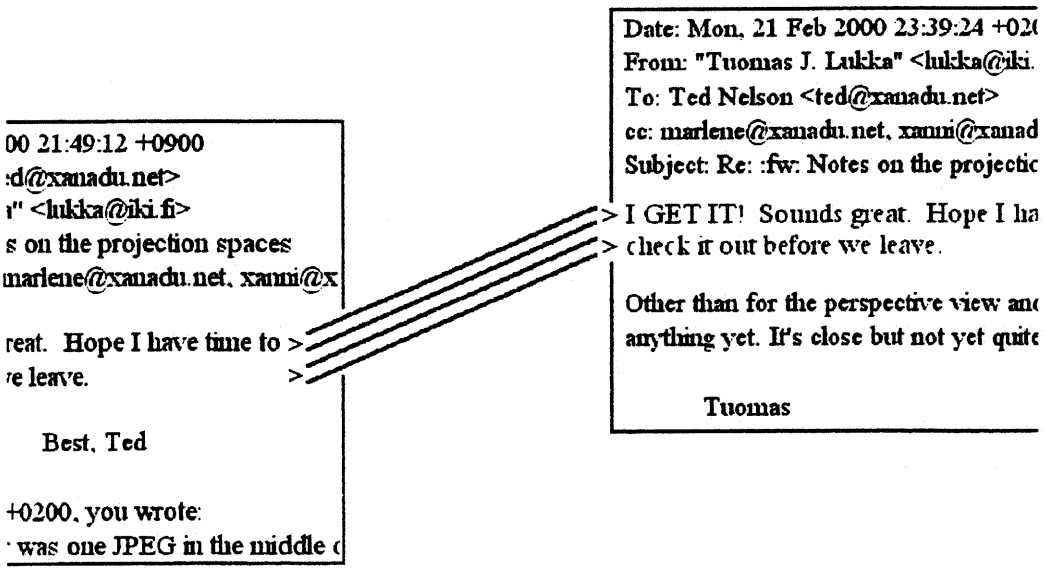

Fig. 21. Simulated email viewer showing original context of a transclusion (rather than the usual disconnected quote). Customary inpointing angle brackets on the right indicate the transcluded material; outpointing on the left, they indicate the source.

Other straightforward examples of where readers might like to have the original context available from every excerpt: interviews and reviews. Another example is programming itself: it would make sense to follow the movement of code visually in big programming projects.

On a small scale, being able to step to the origins of things may not seem important; but when thousands of documents inter-reference one another, the confusion-- and need-- increase. As tomorrow's cosmic Niagara of digital media grows (which is guaranteed not to stay in predictable places), deep context-following, regardless of location of original, will become vital.

MANAGING THE GLOBAL SPACE

There have been various proposals within the Xanadu project for ways to allocate and manage a distributed global address space. (The rest of the world is experiencing the difficulties of this problem now. Internet authorities manage two parallel address spaces centrally: the out-of-sight hierarchy of IP addresses, and the more public and emotive domain-naming system. Both of these are distributed systems with centralized root node management, and the resulting political problems are growing into strange forms.) By contrast, the rootless tumbler addressing system of xu88 (**12**, designed by Gregory and Miller) is particularly interesting for its breadth and parsimony.

Another key question about the universal address space is how to unify it for analysis; especially, within the model presented here, bringing together all the document and link lists to find and compare all the addresses (a serious combinatorial challenge). The Xanadu project pursued two different radical enfilade-based optimization strategies (built into xu88 and xu92). But newer approaches have arisen, such as the search-engine industry, which might perform such services; or the Internet Archive, which could be globally searched for commonalities of both transclusion and content link (**105**).

•

Appendix G
The Transcopyright Concept

To Ted Nelson Home Page

Appendix G.
The Transcopyright Concept.

This is excerpted from **"Xanalogical Structure, Needed Now More than Ever: Parallel Documents, Deep Links to Content, Deep Versioning and Deep Re-Use"**
at http://www.sfc.keio.ac.jp/~ted/XUsurvey/xuDation.html

Theodor Holm Nelson, Project Xanadu* and Keio University

TRANSCOPYRIGHT, PERMISSION TO TRANSCLUDE PUBLICLY

In the system outlined here, a document is basically a list of contents. Thus inclusion of contents in a document is quite independent of ownership of that document, since anyone can list anything.

Then why not allow content to be transcluded freely, by these virtual listings? And allow delivery to be arranged as a separate step? Suppose we separate the issues of who owns the contents, and who may include them?

This leads to a very simple copyright solution. But first let us consider the problem it may solve.

THE TWO PARODY VIEWS OF COPYRIGHT

Copyright thinking has been polarizing into two nasty and crazed views: the "we'll steal it all" school, anticipating the destruction of all copyright and copyright law, and the "we'll nail you for it and lock everything down" school, with pay-per-view methods and client display software that clangs shut when the user's money runs out. These two factions-- copyright hawks and anti-copyright hawks-- are currently engaged in legal and illegal maneuvers throughout the world.

Regrettably, many seem to think these polarized, harsh extremes embody the only possible copyright views.

It is of course conceivable that copyright will be overthrown, as many youngsters hope. But my assumption is that the copyright wars will get nastier and nastier, and the polarization is likely to worsen. As bandwidth and storage increase, the kids will steal and swap more, and crackdowns and unpleasantness will become more frequent and more likely.

The flagrant copyright violators are counting on public sentiment to keep them out of jail. But as with marijuana possesion and sale, which huge numbers of people consider no crime, copyright violators may well start getting prison sentences en masse.

WHAT BETTER ALTERNATIVE?

I would gladly live in a world without copyright, but I don't think that is going to happen. Therefore the best objective, most beneficial to all parties-- and a Xanadu proposal now for forty years-- is to find some way to make copyright less painful, and to facilitate well-intentioned uses of content.

As a rule, finding a principled basis for such a political solution can be very difficult. But by luck, Xanalogical structure directly provides such a principled basis for a win-win copyright method.

A COPYRIGHT APPROACH FOR FRICTIONLESS REPUBLICATION

There exists an alternative, legal copyright system, different from both extremes, which is demonstrably fair to all parties. It is an anomalous system with no close relatives. (We have recently adapted to the Internet the original Xanadu copyright model, formerly centralized (**34**), but now basing it on the same transclusion and payment concept in a new decentralized version.)

Transcopyright (**36, 37**) is a unique alternative copyright system-- an honest system granting users more power, but wholly compatible with existing law and ownership. It takes away nobody's rights. It allows unlimited virtual re-use of contents: all parties may republish, reorganize, anthologize without negotiation, and place excerpted materials into texts, maps, diagrams, time-lines, overlays and so on. And readers may go from every quote to its original context.

Other digital copyright proposals seem principally concerned with sharpening and fortifying a sort of Maginot Line between the two parody views; while transcopyright grants new kinds of re-use to the public-- whose important interest in flexible re-use has not been represented.

The idea is simple. Since the distributed unit is the *content list*, and the delivery of the document is in two phases-- first, the content list; then the content-- why not control the content separately? Why not simply let anyone include any content on their lists, and allow the rightsholder to control the delivery of the content?

We call the legal doctrine *transcopyright*; what it makes possible we call *transpublishing* (**41, 43, 46**), or virtual republication of excerpts without limit from the contents of participating rightsholders.

Transcopyright is a unique bargain, which any rightsholder is free to accept or reject: in return for each excerpt being tethered to its original context (connected transclusively), and therefore no quote being out of context (by inquiry), the rightsholder gives permission in advance for all transclusion in any amount, in published on-line documents, and agrees to furnish any content portions requested. For those crass enough to want to be paid-- and who have some hope of massive downloading of excerpts-- a gateway micropayment may be added.

These two deals represent two levels of compromise and accommodation. The first level of acceptance means you give permission for the re-use in return for the original context being available. The second level means you want the original context available, but you also want to get paid for each downloaded excerpt.

However, the implementation for these two deals is the same at the server level. To get the payment, a gateway micropayment system (such as Hypercoin(tm) **57, 56, 58**) is interposed by the rightsholder.

Many have misunderstood this to mean that transpublishing can only work if a live connection to each unique original is in place; this is a misunderstanding. The same content may be made available from many sources in parallel, provided that they work within the system. Rightsholders are free to operate the servers dishing out their own content, or to delegate that job.

A key factor is that payment should be *proportional* to an excerpt, and small. The gateway micropayment client should have a settable threshold that frees the user from seeing payments less than a chosen amount: "If I click and it costs one cent or less, just buy it."

The example that seems to make the transpublishing idea clearest to people is the following: think of all the movie reviews you've seen on TV where you've wanted to see more-- but not necessarily the whole movie. Imagine movies published on line under transcopyright, from which anyone is free to include favorite scenes virtually-- and any viewer of these new compositions is free to click to see more

beyond the favorite scene, jumping to its context in the original.

Consider that this arrangement turns all participating material into easily-reusable boilerplate, which you can add to any on-line document without discussion. If the idea of microsale seems far-fetched for text, is it far-fetched for streaming video? Or for high-resolution movies? Or the use of unedited footage which is posted for this purpose? What about diagrams? What about historical footage? Surely there will be a niche for this method.

PERMISSION FORMATS

We have developed specific formats for transcopyright permission (**42**) that lawyers consider to be universally legal (since such permission is a license and does not affect existing copyright law). Refining these permission licenses for simplicity and ease of understanding is a continuing endeavor.

INDEPENDENCE OF XANALOGICAL STRUCTURE

The transcopyright method can be used without the linkage part of xanalogical structure, as long as some transclusive delivery method is employed. Indeed, transclusion formats are possible for this using nearly-ordinary HTML (**44**; unfortunately, these formats have had to be tweaked for variations between Netscape and Explorer, and the results look somewhat different in each). At Keio University we have produced an experimental server and editor for transpublishing with these formats (**46**), under a grant from the Japanese government.

CACHING, SPEED AND RELIABILITY

It is true that transpublishing requires some kind of network, but it not necessarily perfect connectivity; it should be possible for transpublishing to work efficiently on imperfect networks through caching systems designed for the purpose.

- "But that means if the network goes down I can't get it"-- true, but that's the same as for any other downloads.

- "It's too complicated." As compared to what other system that allows you to include copyrighted material?

- "It will all get stolen." Right. Unemployed napsters and napstresses are going to cache millions of hours of video. And then what will they do with it?

ACM'S INDIRECT ENDORSEMENT OF TRANSCOPYRIGHT

The transcopyright system has been endorsed in principle by the Association for Computing Machinery in its current "Interim Copyright Policy" by Peter J. Denning, presently at the ACM site (**38**)*. To quote:

> "Transcopyright permission for electronic dissemination. ACM incorporates a principle similar to one Ted Nelson called "transcopyright". ACM will hold its copyrighted works on its servers and will give free and unlimited permission to create and copy links to those works or their components. So that readers can locate the context from which an excerpt was drawn, ACM will provide a way of linking a component to its parent work."

Laudably, this is intended to maintain connection to the original context, as few Web links do. Unfortunately, the ACM statement is silent on the central tenet of transcopyright, which is the permission freely to recomposite and republish virtually by some transclusion mechanism. Still, the ACM's approval-in-principle is very heartening.

SUMMARY OF TRANSCOPYRIGHT

Transcopyright is a copyright system which is open, liberal, benign, fair, frictionless, honest, cheap, and win-win. It is intended to facilitate ease of access to partial documents; to facilitate ease of recomposition and republishing as if contents were in the public domain; and to do what paper cannot: provide access to the original context of any excerpt.

It was always a premise of the Xanadu project that in the future people are going to read and view fewer and fewer finished works as the avalanche of media increases; and so paying for only what you want (and getting to keep it), but knowing you can get the rest, is what we really need.

•

Appendix H
1999 HTS Design

tpubstrux
99.08.01 (d28r

To Ted Nelson Home Page
To Transpublishing Main Page

For FURTHER INTRODUCTION, see our other writings on the subject–
• further introduction to the idea of transpublishing;
• the theory and rights aspects of transpublishing;
• the benefits of the system;
and
• the type of micropayment mechanism which may be added later.
FURTHER TECHNICALITIES to be found elsewhere include:
• Keio formats for transpublishing
• the vocabulary necessary for this project (a listing)

Keio-JIPDEC Hypertransaction Project

TRANSPUBLISHING STRUCTURES FOR TODAY'S WEB

Theodor Holm Nelson
Keio University Shonan Fujisawa Campus
and
University of Southampton

"Literature" may best be understood as "the media we keep". Hypertext is the natural extension of text media, and so the design of hypertext systems determines the future directions of much of literature. What follows is a technical design intended to achieve literary purposes, such as the connection of on-line textual quotations to their original contexts. These proposed specifications were developed as part of the HyperTransaction Project at Keio SFC, as a step toward implementing an improved

quotation system for the World Wide Web.

The World Wide Web in its present form is based on a simplistic model of connection, composed only of separate pages which change separately without coordination and have links that point in only one direction. This system has numerous drawbacks. Three of them are:

- There is no way that a quotation can connect to its origin, so the reader may see the original context;
- there is no way that material may be republished without all the same negotiations and permission necessary for paper publishing;
- there is no way for authors or publishers to be paid for re-use of their materials.

We want to improve this in several ways with a very simple idea: transpublishing. Transpublishing means bringing into an electronic document quoted contents from another electronic document. Transpublishing, in the form we foresee, will mean being able to republish copyrighted material on line, in any quantity, without negotiation, to be delivered each time from the original publisher.

Transpublishing derives from Project Xanadu (1, 2), as more recently adapted to the Web with the transcopyright legal system (3, 4).

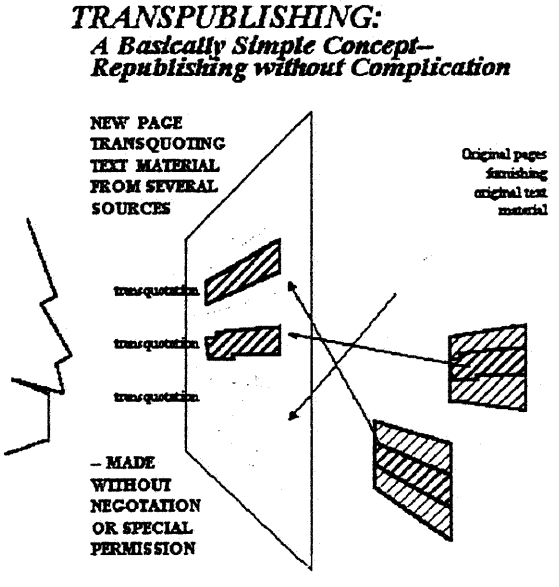

Conceptually, transpublishing is done by invitation: the HTML in a Web page invites various sources to deliver quoted materials. These invitations are sent out by the browser. When the materials arrive, the browser assembles them into the visible page. Instead of placing the actual text of a quotation on a page, an author wishing to include a transquotation places a particular HTML string in the source code of the

desired Web page. This invites the delivery of the quotation from the original, and places it in position in the finished page.

Legally, in order to be generally useful, transpublishing requires an arrangement that permits others to include material virtually without negotiation. A legal permission basis for this has been established and approved by lawyers (3).

Technically, transpublishing is the process of virtually republishing content which is already published by means of "inclusion by reference"-- that is, not by copying the material into the newer document, but using mechanisms which cause the republished portion to be downloaded from the original (or some logical equivalent).

Particular advantages of transpublishing include:

- FREEDOM TO RE-USE: Any materials within this world may be freely recomposited without negotiation-- that is, put into any new on-line document in any quantity without contacting the copyright rightsholder, much as if they were in the public domain.

- CONNECTION BETWEEN QUOTATION AND ITS ORIGIN: Any quotation is connected to its origin.

- MICROPAYMENT TO ORIGINATOR: A payment system can be added whereby each recipient pays for each separate quotation. This would mean that rightsholders may be paid for their Web publications in exact proportion to usage.

The general objective of the Keio Hypertransaction Project is to take a first step in this direction, creating software for the Internet for connecting quotations on a Web page to their original contents, and allowing everyone easily to quote materials available under transcopyright. The particular objective of the Keio Hypertransaction Project is to create server software for individuals, organizations or governments to facilitate publication of on-line text which may be quoted with Web connection to its original context.

Transpublishing Formats

We have developed formats and procedures to simplify and clarify this process, designed to deliver quotations smoothly to ordinary browsers while maintaining the connection of the quotation to its original context. These formats have visual components, connection components and HTML structure underneath.

Visually, the Keio formats have special copyright and quotation marks (TCOmark and TQmarks). These marks connect a picture or quotation to a publisher's permission statement which allows this usage, and connects a text quotation to its original context.

PHILOSOPHY OF HYPERTEXT

The Keio formats are different for text and for pictures. The formats include:

- **permission statements**, so that a publisher may give easily-found permission allowing contents to be transquoted;
- **marks** to indicate that contents that such permission exists;
- **exact forms of connection** between such transquotable contents and the permission statement, so that other publishers may check the exact details permitting quotation from that original document;
- **TQstrings**, particular HTML to be inserted by a transquoting party, which invite the server to send a transquotation and maintain the connection between the transquotation and permission statement.

A transquoting page may have six different types of parts:

- New material belonging to the current author, which is not transquoted;
- Transquoted picture (also called transpicture or transpic);
- TCOmark, a transcopyright mark for a transpicture (connected to the transpicture's permission statement);
- Left TQmark, the special left quotation mark of a transquotation (connected to the original document from which a transquotation is taken);
- transquoted text material (also called transquotation or transquote);
- Right TQmark, the special right quotation mark of a transquotation (connected to the permission statement of a transquotation).

These are visually quite recognizable.

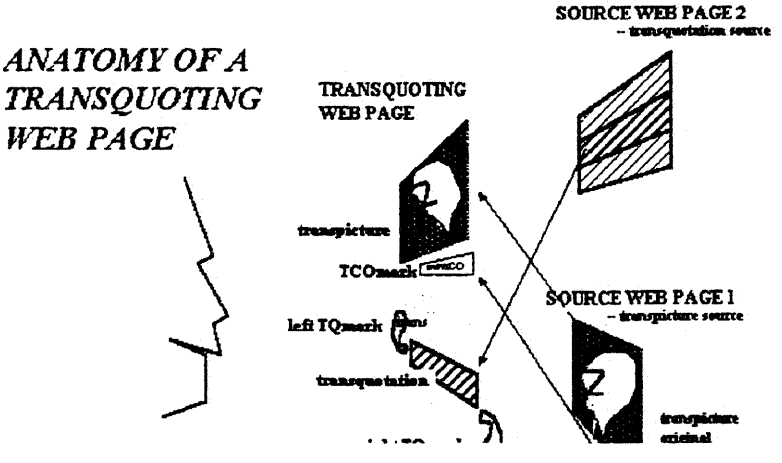

PHILOSOPHY OF HYPERTEXT

USER

TCOmark

In the Keio transpublishing formats as we have developed them, specific HTML strings bring in these marks and maintain all these connections. Such HTML strings we call TQstrings. A TQstring is the sequence of HTML which brings in a transquote or transpic, including the special copyright marks, and keeps those marks connected to the permission statement and (for text) the original context page.

TQstring for pictures

A transpicture (transpic) does not need any new kind of tag, since ordinary HTML (the IMG tag) will transpublish it from anywhere, putting it correctly into place on the final page. To make such transpublishing of a picture legal under transcopyright, it needs only a transcopyright mark (TCOmark) to link it to the permission page.

Using this format, a transoriginator indicates availability of the picture, and a user may find further details in the permission statement.

Thus the string consists of an IMG tag for the picture, an IMG tag for the TCOmark, and the link surrounding the TCOmark, which connects to the permission statement.

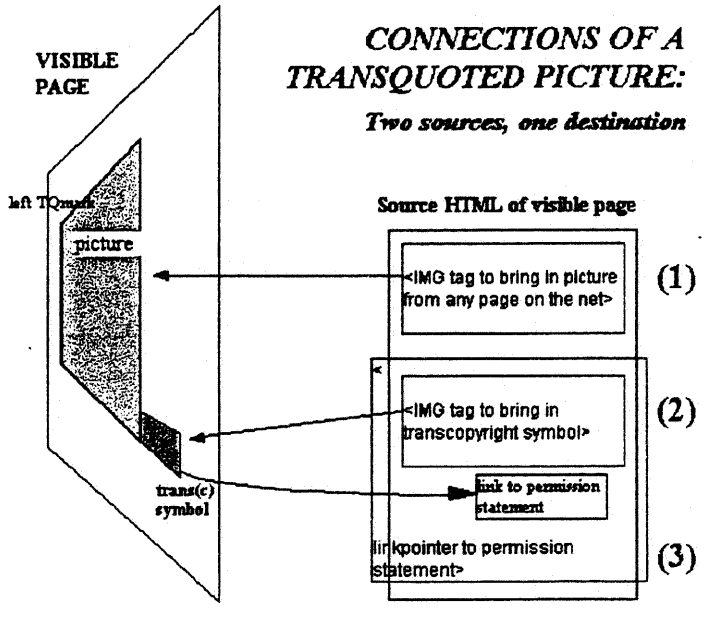

Sample:

```
<img SRC="SasagawaSelfportrait.gif" BORDER=0 height=155
width=123><a

href="http://www.sfc.keio.ac.jp/~ted/TPUB/SasagawaPermish.html"><

    SRC="tcosymb.gif" height=14 width=36></a>
```

This brings in both a picture and its associated TCOmark, connected to its permission statement, as follows (picture by Keio student Kouhei Sasagawa):

The TCOmark functions like an ordinary copyright mark (which it contains), informing the viewer that copyright is claimed on the work, but additionally provides a link to the legal permission which allows it to be recomposited into any on-line document, such as a Web page, without contacting the rightsholder.

This method requires no additional technical implementation except editing software to make it more convenient for authors to include the picture and its connected TCOmark.

TQstring for text

A textual transquotation, in the Keio formats, looks like this (simulated):

Correct Transquotation View

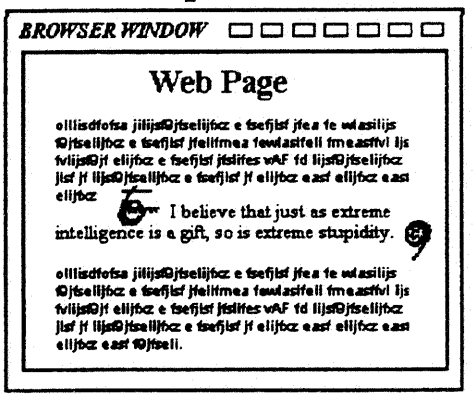

The left special quotation mark (TQmark-left) connects to the original context of the transquotation (the original document), and the right special quotation mark (TQmark-right) connects to the permission statement for the original document.

This format offers

- a transcopyright symbol to show that the material is owned but available under permission;

- a path to the permission page for the original document;

• a connection to the original context of the text quotation.

• a possible error indication: if no contents appear between the transquotation marks, their presence with no contents indicates a malfunction.

Let us look at the HTML source for these parts of a transquotation. It is made up of the TQstring-left for the left TQmark (linked to the context page), the TQtag, and the TQstring-right for the right TQmark (linked to the permission page).

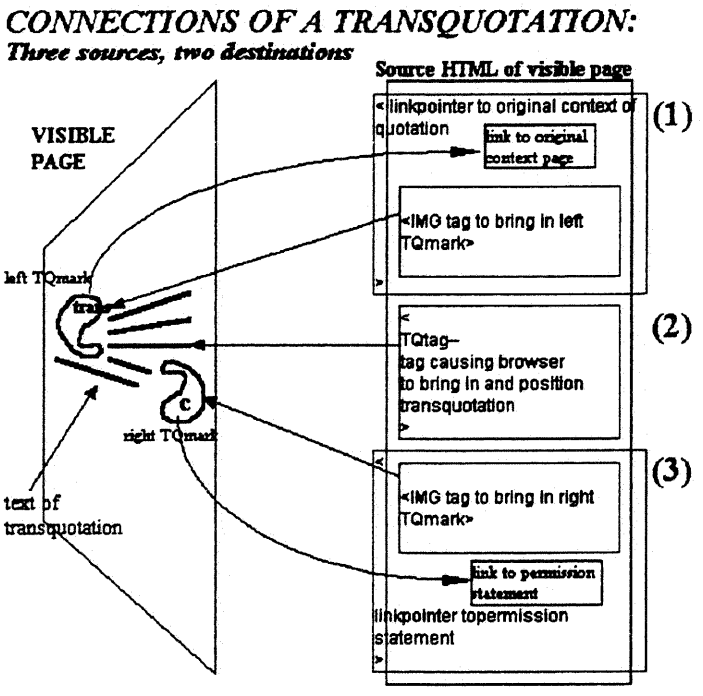

Section (1) of the TQstring for text brings in the left TQmark and connects it to the original document. Section (2) is intended to call the quoted material itself and position it in the Web page. Section (3) brings in the right TQmark and connects it to the permission statement.

The only part of this which does not yet work under Web-standard methods is part (2), the transquotation tag, which we call the TQtag. There is no existing standard format to specify a quotation to be carved out of a document, and no existing standard format to place it within text as an in-line sequence. Our present interpretation of the current status of HTML and XML requires that we substitute a temporary mechanism which can be exactly replaced when a satisfactory tag has been standardized.

Such a tag will necessarily contain the URL for the quotation, and embed the quotation brought in by that URL into the text, surrounded by the TQmarks we have designed. Within the syntax customary in HTML, such a TQtag will have a beginning (such as "<TQtag>") and a matching end (such as "</TQtag>"), although we do not know exactly what they will be.

Interim Proposal for TQtag Functionality

Our interim proposal for handling the functions of part (2), the missing TQtag, is in four parts. Such methods will have to be used until some method (perhaps such as ours) becomes standard.

(2a) to request the transquotation with a quotation-extraction format

URL/docname#[n+m]

where "URL" is the URL of the page to be quoted, n is the character position of the beginning of the quotation and m is the length of the quotation;

(2b) to create a webserver (the TQserver) which will respond to this request;

(2c) to find some convenient temporary mechanism in today's HTML for positioning the quotation in an output document;

(2d) to provide (for ideal viewing) a proxy server that creates the desired appearance on the final page.

Sample TQstring for a text quotation:

```
<A HREF="http://www.sfc.keio.ac.jp/~ted/TPUB/Superman.html">
<IMG SRC="aposleft.gif" BORDER=0 HEIGHT=40 WIDTH=37></A>
<TQTAG> <--TQtag firstpart-->
"http://TQserver1.sfc.keio.ac.jp/Superman.html#[29+80]"
<--TQtag lastpart--> </TQTAG>
<A
HREF="http://www.sfc.keio.ac.jp/~ted/TPUB/SupermanPermish.html">

<IMG SRC="aposrt.gif" BORDER=0 HEIGHT=49 WIDTH=43></A>
```

The tag opening and closing <TQTAG> and </TQTAG> represent whatever tag mechanism will eventually be used in a standardized system. The segments in this example to be replaced by some convenient meechanism of choice are the two strings <--TQtag firstpart--> and <--TQtag lastpart-->.

The simulated result was shown above ("Correct Transquotation View"), extracting the transquotation from the document "Superman" by Keio student Rica Miyata (5).

A LARGE-SCALE SYSTEM

PHILOSOPHY OF HYPERTEXT

We envision a large-scale transquotation system which could function world wide as a part of the Web. It would have four parts:

1. Transquotable Server (TQserver)

This server delivers a quotation as requested, extracting it from a larger document and returning it to the browser.

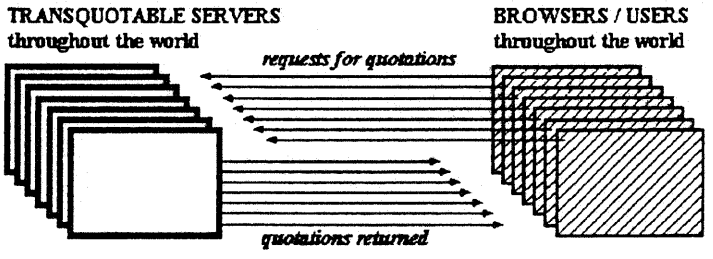

2. TQRegistry

In order to publish a transquotable document, a person or organization would submit the document to the TQserver in an appropriate form. The part which receives the document is called the TQregistry. A TQregistry must be associated with each TQserver. The TQregistry puts the necessary information for each document into a database, used by the TQserver to perform its tasks.

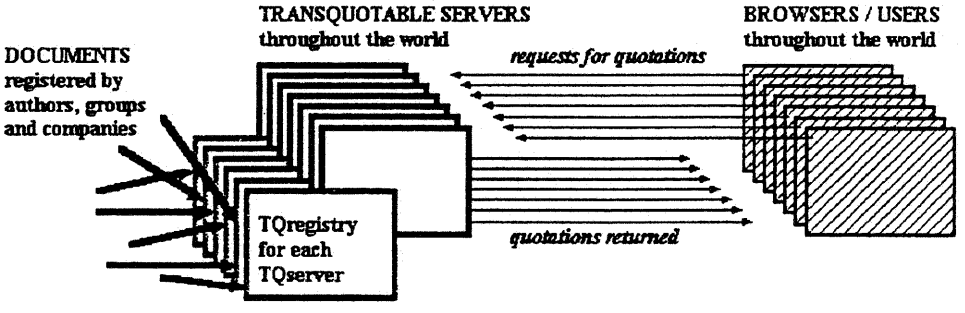

(3. Micropayment System)

The eventual design provides for a possible micropayment system, to be used by some publishers at their discretion. That is not part of the current project.

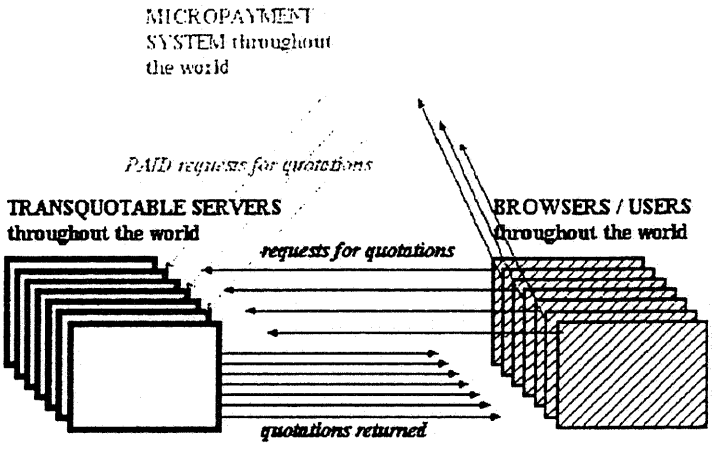

4. The TQproxy

For the present time, it will also be necessary to offer a Web proxy server to give the pages their correct appearance. This is due to a problem with current HTML tags, which we expect will be corrected shortly. The TQproxy will not be necessary in the final system.

Final Note on Caching and Virtual Delivery

It is clear that Web caching, as usually implemented, would conflict with the this system's intention of downloading small portions from individual publisher.

However, some form of caching which does the equivalent-- effectively delivering or selling content from the separate originals-- should be technically possible. We consider this an optimization issue to be addressed later.

REFERENCES

1. Theodor H. Nelson, *Literary Machines*, 1981. Most recent edition *Literary Machines 93.1*, 1993. Mindful Press, Sausalito, California; distributed by Eastgate Systems, Cambridge MA, USA. Translated also into Japanese (ASCII, Tokyo, ISBN 4-7651-0418-5.)

2. Theodor H. Nelson, "The Heart of Connection: Hypermedia Unified by Transclusion." *CACM* August 1995, vol 38 no 8, pp. 31-33.

3. Theodor H. Nelson, "Transcopyright: Dealing with the Dilemma of Digital Copyright." *Educom Review* 32:1 (January/February 1997), 32-5.

4. Theodor H. Nelson, "Transcopyright: A simple legal arrangement for sharing, re-use and republication of copyrighted material on the Net." Keynote at WWCA '97 conference, March, 1997, Tsukuba, Japan; published in Takashi Masuda, Yoshifumi Masunaga and Michiharu Tsukamoto (Eds.), *Worldwide Computing and Its Applications*. Springer-Verlag, Berlin, 1997. ISBN 3-540-63343-X. Pp. 7-14.

5. Rica Miyata, "Superman", at http://www.sfc.keio.ac.jp/~ted/TPUB/Superman.html.

Appendix I
Putting Xanadu on the Web

PHILOSOPHY OF HYPERTEXT

xuModelPage
01.07.14 (d17

To Project Xanadu main page

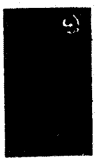

DEEP HYPERTEXT: THE XANADU® MODEL

(After reading this, check out how we're moving it to the Web.)

Today's one-way hypertext-- the World Wide Web-- is far too shallow. The Xanadu project foresaw world-wide hypertext decades ago, and endeavored to create a much deeper system. The Web, however, took over with a very shallow structure. Our simple method, very different, allows--

- UNBREAKABLE LINKS.
- COPYRIGHT SIMPLIFICATION AND SOFTENING: by special permissions and methods, quotations of any size may be used by anyone and mixed together frictionlessly.
- ORIGIN CONNECTION: All quotations and excerpts stay connected to their original.
- TWO-WAY LINKS: anyone may publish connected comments to any page.
- SIDE-BY-SIDE INTERCOMPARISON OF CONNECTED DOCUMENTS-- showing two-way links, differences between versions, origins of contexts. (For a simple working demonstration, see our new free CosmicBook (tm) reader.)
- DEEP VERSION MANAGEMENT: documents may be changed incrementally (with each version available); versions may branch; authors may easily see exact differences between versions.
- INCREMENTAL PUBLISHING: new changes may be continually made by authors without breaking links.

How can this be? Very simple, but very different.

THE XANADU MODEL

The Xanadu model has always been very simple: make content available with certain permissions; then distribute and maintain documents simply as lists of these contents, to be filled in by the browser (in the same way that browsers now fill in GIFs). This list of contents is effectively a virtual file of contents to be sent for and how to put them together.

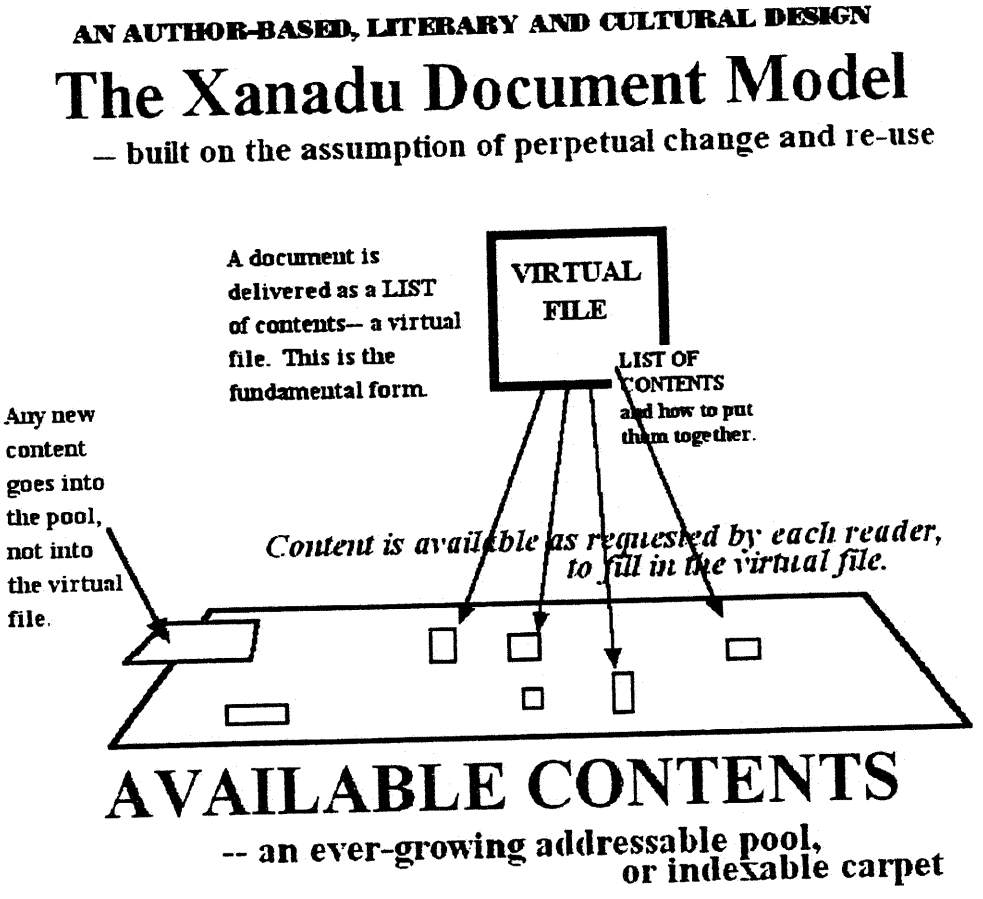

This means

- everyone can re-use all content virtually, simply by listing the desired content.
- Since links are between the addresses of these contents, links are intrinsically bidirectional, may be made by everyone, and may overlap in vast numbers.

- Since the original address of content is known, its original context may be obtained.

We will not go into further details here. These matters are written up at length in various places, e.g. the recent Xanadu technical summary for the ACM.

The Xanadu copyright model-- now called "transcopyright"-- takes a little more explaining.

MOVING THE XANADU MODEL TO THE WEB

We are now moving the Xanadu model to the Web, which is an uphill fight, but we believe it's worth it.

The political fight against the weight of today's standards and methods is enormous. But what we are trying to do is clear up the problems that these standards, and the present way of thinking, have created.

To Project Xanadu main page

THE NEW XANADU® STRUCTURE FOR THE WEB

Today's one-way hypertext-- the World Wide Web-- is far too shallow. The Xanadu project foresaw world-wide hypertext and has always endeavored to create a much deeper system. The Web, however, took over with a very shallow structure. Our simple, but very different structure-- for details see "The Xanadu Model"-- allows--

- UNBREAKABLE LINKS.
- COPYRIGHT SIMPLIFICATION AND SOFTENING.
- ORIGIN CONNECTION.
- TWO-WAY LINKS.
- SIDE-SIDE INTERCOMPARISON.
- DEEP VERSION MANAGEMENT.
- INCREMENTAL PUBLISHING.

How can this be? Very simple, but very different.

The Xanadu model has always been very simple: make content available with certain permissions; then distribute and maintain documents simply as lists of these contents, to be filled in by the browser (in the same way that browsers now fill in GIFs).

Since the advent of the Web, our last several years have been concerned with figuring out how to move these concepts to the very different standards environment that the Web has imposed.

One new design, underway in collaboration with David Durand, is now being written up by him. It is unrelated to the previous Xanadu code. (Further details will be presented at the HT01 Hypertext Conference in Denmark,

August 2001.)

The new design consists of:

- A new file type for virtual content (tentatively called .VLIT, Virtual Literary Format)
- Servers delivering portions on request (using existing protocols)
- Editors for the .VLIT file
- (Later) Browser plug-ins for intercomparing .VLIT files, showing their differences side by side.

The Xanadu copyright model-- now called "transcopyright"-- takes a little more explaining.

The political fight against the weight of today's standards and methods is enormous. But what we are trying to do is clear up the problems that these standards, and the present way of thinking, have created.

To Project Xanadu main page

TRANSCOPYRIGHT FOR THE WEB

A clean legal and technical system for re-using on-line content

The Web has finally discovered the copyright issue, but still at a naive level.

The excitement over Napster has been about the right to download and save content in whole files and chunks. This does not even begin to address the much deeper issue of how we can sensibly re-use content among on-line documents.

For scholarship, for understanding, and for appreciation, we need to be able to quote things in large quantity, in open republication-- but no one has seen the obvious path, which has always been the Xanadu model.

Never mind about the Napster-style thieving by night. The Xanadu project proposes an honest method that can be used in the open by people with day jobs and supported openly by ethical service providers.

NOT ALL CONTENT WILL BE FREE

Lots of content will be free. But lots of other content will cost money; many publishers are seeing to that, and the issue is how to live with these relationships cleanly.

How can we live and work in this real world? How can we be allowed to make detailed and connected comments? (For a sample of how connected comments should work, try our free CosmicBook(tm) reader.)

HOW WE CAN RE-USE AND MIX CONTENTS

Here is the fundamental objective of the transcopyright system:

> TO BE ABLE TO MIX CONTENTS IN ANY QUANTITY FROM ALL SOURCES, FOR COMMENT AND UNDERSTANDING, WHETHER THOSE SOURCES COST MONEY OR NOT.

Many people complain about that fact that this is not now possible, but want the laws to go away, which they won't. We believe we have a uniquely beneficial method.

THE XANADU MODEL

The Xanadu model has always been very simple: make content available with certain permissions; then distribute and maintain documents simply as lists of these contents. The browser then obtains the contents separately.

This means (in principle) that everyone can re-use all contents. The problem is in the details; what follows is a clean and consistent set of details.

THE XANADU MODEL AND TRANSCOPYRIGHT

Here's how this can work on the Web with our new Web formats. (It may sound complex, but it is no more complex than the World Wide Web today-- just very different.)

The transcopyright system is this. As in the Xanadu model:

- a virtual format, listing the contents;
- an editing system that manages this virtual format;
- proportional delivery on request;

And in addition:

- a permission system by rightsholders in advance (transcopyright permission-- see original article and as published);
- for those contents which cost money, a

proportional micropayment system (overlay or gateway micropayment);
- a caching method for content which has already been acquired (Deep Cache).

The following picture ties it all together. It may look complex, but that's because we don't have time to make a simpler diagram :)

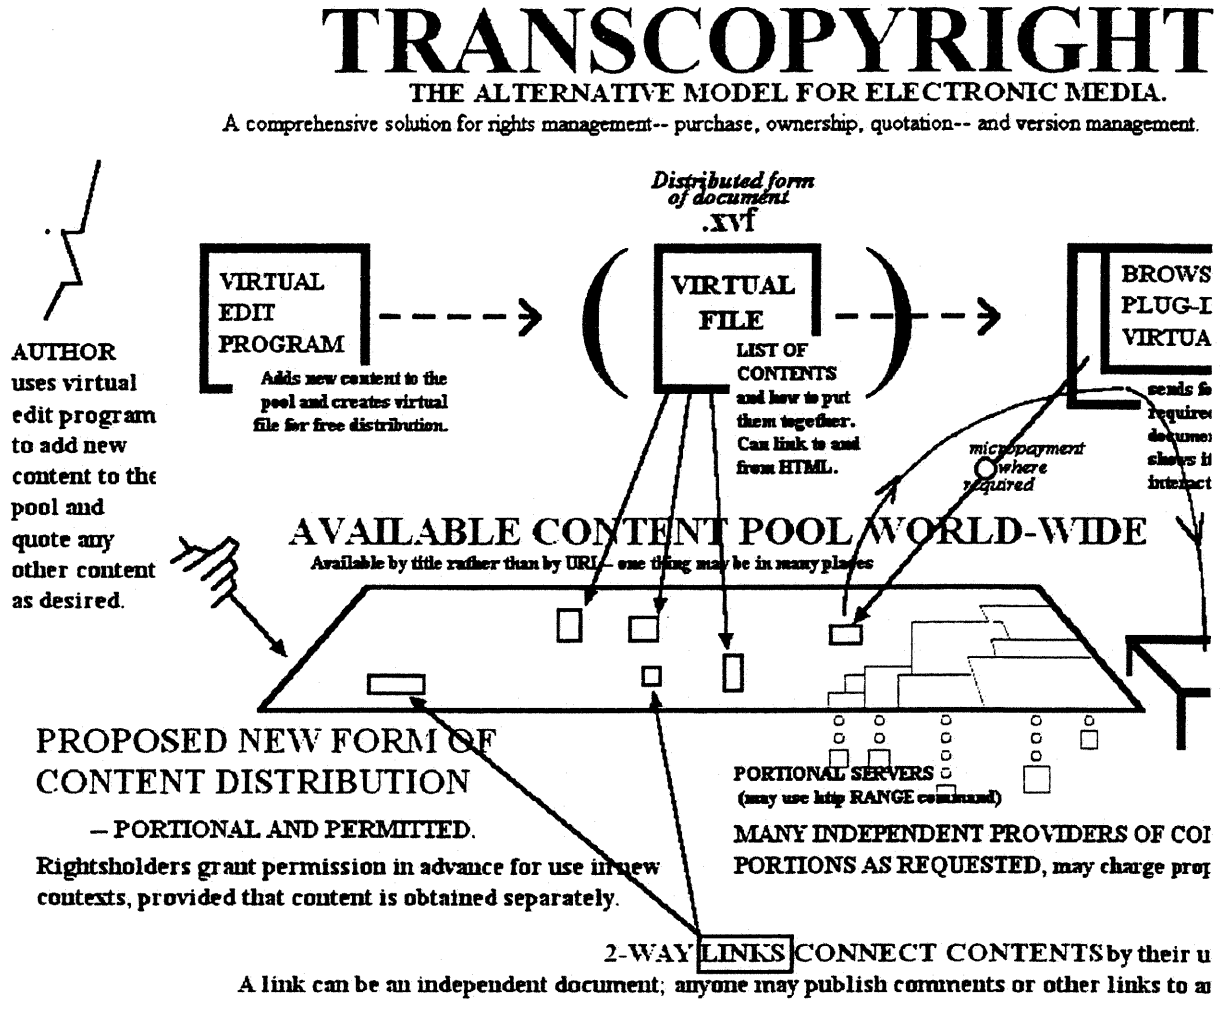

vlit
01.10.16 (d 10z

To Ted Nelson Home Page

The Virtual Literary Format

Ted Nelson

Thanks especially to David Durand of the XML Linking Committee,
and to Tim Brailsford and Craig Stewart of the University of Nottingham,
for their consultation and design assistance.

The Virtual Literary Format

The Virtual Literary Format (VLIT) is intended to be a Web-compatible document format that follows the Xanadu® model, solving a number of legal and technical issues--

- ownership and copyright
- re-use and republication
- nonbreaking links
- multi-ended and typed lists
- intercomparison and versioning
- annotation and controversy
- link origins and management
- archiving.

These problems are handled in a unified and minimalist fashion, as long intended by the historical Xanadu designs, rather than in separate, specialized and tangled methods like those proposed amongst the Web community.

The general concepts may be found in the historical Xanadu article. We will discuss here the general concepts of the format, then difficulties of implementation in the present Web environment.

Structure of the VLIT file

The VLIT file contains no data, but brings in text, audio and video from stabilized long-term sources on the Net. (It is essentially the same as a video EDL file, generalized.)

The VLIT file is to be in two sections:

- CONTENT LIST: pointers to content to be brought in from stabilized long-term sources. The content will be mainly series (or spans) of discrete content elements (characters, audio samples,

video frames) to bring in. The countable sequence of these elements to be the central reference method of the file.

• PRESENTATION OVERLAY: directions for presentation ("markup") are kept in a parallel overlay section of the file. This means that the contents can be presented in new ways even though the original content cannot be changed, and regardless of the format of that original content. These markup spans reference the numerical positions of the elements in the content list.

• LINKS, if any, are addressable entities which appear in a third list and reference the stabilized universal content address.

THE LOSS OF PRIOR STRUCTURE

A principal concern is to allow new uses of content. In order to allow restructuring in new uses, prior structure is not recognized in the VLIT file. (Showing old structure to the new author should take place in the VLIT editing program, or to LATER later readers through intercomparison windows.)

The EDITING PROGRAM

The VLIT format must be edited through an editing program which fulfills the content, then strips it for distribution.

"WEB COMPATIBILTY"

"Web compatibility" is a slippery concept. The World Wide Web Consortium wants to be the controlling organization for this concept.

However, since the World Wide Web Consortium has a completely different idea of the meaning and purpose of hypertext, their standards are generally irrelevant to this enterprise.

In any case, as far as the public is concerned, the Web is anything that comes from the Internet into a browser window.

Thus the real problem for our project is to make the VLIT contents open in the standard browsers-- specifically, Explorer, Netscape, Opera, Amaya and Mozilla.

PLUG-IN VERSUS RAW OPENING

Since the VLIT diverges from the standard Web approach, it would be easiest to do this with our own browser plug-in. This allows for the greatest flexibility and independence.

However, such a plug-in could reduce user acceptance.

If possible, it would be desirable to show these documents in an ordinary browser window without a plug-in. While this is supposed to be possible according to W3C standards, there are problems.

TO: PROBLEMS OF VLIT

•

To Ted Nelson Home Page

Problems of the VLIT Format

Ted Nelson

Thanks especially to David Durand of the XML Linking Committee,
and to Tim Brailsford and Craig Stewart of the University of Nottingham,
for their consultation and design assistance.

Problems of the Virtual Literary Format

Continued from "The Virtual Literary Format"
(separated to make updating easier,
since the problems will probably change
more often than the basic idea therein.).

BACKGROUND:

The Virtual Literary Format (VLIT) is intended to be a Web-compatible document format solving a number of legal and technical issues, including--

- ownership and copyright
- re-use and republication
- intercomparison and versioning
- link origins and management.

The VLIT file is to be in two sections:

- a list of discrete content elements to bring in, in countable sequence of these elements;
- a parallel overlay section whose addresses reference the numerical positions of the elements brought in

---

THE PROBLEMS

It turns out that using W3C-recommended methods is currently extremely difficult-- for practical reasons having to do with the stability and implementation of those advertised standards.

The World Wide Web Committee has proposed a wide variety of XML-based methods of referring to things, with overlapping capabilities for pointing and connecting. Unfortunately these standards are in various stages of development and deployment--

- not all of them are finalized as specifications;
- even for those specifications which are finalized, they are differently implemented, incompletely, on different browsers, including the W3C's 'Amaya' flagship browser.

MOST APPROPRIATE METHODS, IF THEY WORK

My advisers agree: the two W3C methods which should do the job most cleanly are XINCLUDE (for the content list) and XPOINTER (for the presentational markup). The details need not concern us here, as they should be should be readily apparent to anyone who knows these methods.

Unfortunately, tests by my advisers indicate that none of these browsers (Explorer, Netscape, Opera, Amaya and Mozilla) to be sufficiently compliant with the W3C standards for XINCLUDE and XPOINTER to use them for implementing VLIT at the present time.

TACTICAL ISSUES

For our intention to implement the virtual literary format, there are several options--

1) WAIT-- until the official bodies and implementing companies to converge on standardization.

- Problem: We have no assurance that they will converge on standard implementation in the next five years, if ever. (Consider the continuing variations of JavaScript between Netscape and Explorer.)

2) SIMULATION BY JAVASCRIPT AND RELATED METHODS. A solution proposed severally by Durand, Brailsford and Stewart is that we implement in the file format some combination of code (especially Javascript) to simulate the intended effect. (We have done this in the Keio Transpublishing Project.0)

- Problem: Such a format obfuscates and entangles the simplicity of the VLIT design.
- Problem: This is subject to varieties of Javascript

implementations, and their possible changes.

• Problem: This may only be an approximation, with inconsistent and undesirable features.

• Problem: This format might then become inferior and obsolete, if and when standardization converges.

3) PLUG-IN FOR A WHOLLY DIFFERENT FORMAT. There is a case for creating a fully alternative format and a plug-in to present it, programmed in anything from Python to ZigZag. While this is a further divergence from W3C's intended standards, it also releases us from many concerns which are both irrelevant to our cause and could waste resources.

• Problem: the psychological issue of noncompliance to W3C "standards" could work against us.

However, that has not prevented RealNetworks or Flash from becoming accepted on the Web.

Moreover, support from a major source (such as the French government) could overcome such resistance.

CONCLUSION: THE CHOICES IN BRIEF

We should do whatever is necessary to make the system work as soon as possible. Waiting for others is rarely the productive way to proceed.

Even to use the W3C's chosen methods, we will have to create code to implement the W3C methods correctly in order for the VLIT methods to work.

If, however, we need to do serious programming, is there any point in simulating the irrelevant structures of the W3C while we are waiting for them to come into synchronization? Or should we build a better format and plugins for it?

Appendix J
Xanadu Design Ideas

ApJ- xuDesignIdeas
01.11.04 (d24y

To Ted Nelson Home Page

Appendix J.
Xanadu® Design Ideas

This is excerpted from "**Xanalogical Structure, Needed Now More than Ever: Parallel Documents, Deep Links to Content, Deep Versioning and Deep Re-Use**"
at http://www.sfc.keio.ac.jp/~ted/XUsurvey/xuDation.html

Theodor Holm Nelson, Project Xanadu* and Keio University

Xanadu Design Ideas

Unusual design ideas of the Xanadu Project have included (keyed to the standardized list of Xanadu implementations, with best recollection as to dates), and no attempt to sort out credit.

We will also not consider here those other parties who are credited with some of these ideas, or who was first. These are interesting and important issues, and of great importance for professional courtesy, but much too tangled for the present discussion.

- hypertext concept, xu60 (term introduced 1965)
- hypertext links, xu60
- referential editing, xu60
- transclusion, xu60; item-level transclusion with parallel display, xu65; character-level transclusion in a universal address space, xu88
- branching versions, 1960
- copyright solution by referential transclusion, xu60
- network storage for documents paid for by publisher, xu60

• transpointing windows, 1965 (term introduced 1994)

• proposal for email, 1965 (implicitly proposed in illustration)

• parallel markup, xu70

• ring buffers for text, parallel ring buffers for parallel markup and editing, multidimensional ring buffers, xu70

• threading a display processor through the main data structure (rather than generating a separate display list), xu70 and xu99

• event loop driven by user input (JOT™, xu72)

• piece tables, xu72 or earlier

• enfilades: the model T enfilade (a data management tree for indexing linear text and editing it by reference), xu72; tightly coupled enfilades, xu76 and xu88; two-dimensional enfilades, xu88

• general enfilade theory, ca. 1980 (generating data management trees with an upwardly propagating search property and simultaneously a downwardly imposable structural property)

• fluid media with stabilized universal addresses, xu88 (term "fluid media" introduced 1999)

• generalized xanalogical protocol based on lists of content spans, xu88

• tumbler addressing (multi-segment universal addresses), xu88

• tumbler span arithmetic based on transfinite arithmetic, xu88

• tumbler span matrix operations, xu88

• version management by tumbler span permutation matrices, xu88

• search of entire docuverse for overlapping spans by 2D enfilade, xu88

• the ent (a singular enfiladic structure with built-in versioning), xu92

• multidimensional list structures with cursor-centric viewing in rows and columns zz86 (ZigZag(tm) structure), zz86

- transcopyright doctrine separate from software, transco95

- transpublishing within conventional HTML, 1996

- transclusion tag and protocol specifications to work in ordinary Web browsers from special transpublishing servers, hts99

- N 1/2D graphics, xu99

- multidimensional computer structures and graphics with a ZigZag spine, xu99

- ND mapped to 3 1/2D for high-performance use of 3D graphics boards, with OpenGL threading through 3D views of the ND structure, xu99

- profuse interpenetrating structures as an alternative to "applications", xu99

- high-performance 3 1/2D as a potential alternative to the "desktop", xu99

- profuse interpenetrating structures as an alternative to "applications", xu99

-

-

Appendix K
Principal Xanadu Internal Designs

To Ted Nelson Home Page

Principal Xanadu Internal Designs

These are the principal designs for Xanadu internal mechanisms that are known to me. (However, other members of the project may remember other significant designs.) For the sake of discussion here, I have given coordinated nicknames to these principal designs. Unfortunately documentation is spotty and being attempted retroactively. At least ten of these designs (marked with asterisks below) have been implemented to some state of demonstration or delivery. Dates selected are those of significant closure.

• xu60-- 1960 preliminary design (currently undocumented)

* xu65--1965 zipper lists (23); implemented 1996 by Kiyoki Ookubo and others, using Yuzuru Tanaka's IntelligentPad system (115)

• xu67-- 1966 SNP design (Sexus-Nexus-Plexus) separated data into content, referential lists and links (currently undocumented)

• xu70-- 1970 parallel markup, ring buffers (27)

* xu72-- 1972 first enfilade (model T), enfiladic referential editing using word-processor interface (28)

• xu76-- 1973-6 designs by William Barus (currently undocumented)

* ent80-- core design of "ent" versioning enfilade system by K.Eric Drexler (26)

* zz86-- multidimensional list system with cursor-centric views; now called ZigZag(tm) (71, 70); prototype implemented 1997 by Andrew Pam and variously downloadable (73), instructions at (75); extensible with explorable genealogy demo and genealogy demo instructions (77); Java version now being developed under open source ("Gzigzag") at Sourceforge.net under stewardship of Tuomas J. Lukka (recent version downloadable (26)).

* xu88 (formerly called Xanadu 88.1 and now called "Udanax Green")-- distributed client-server system programmed in C,

principally designed and implemented by Roger Gregory and Mark S. Miller, partly based on the general enfilade theory of Mark S. Miller, Stuart Greene and Roger Gregory; for fluid media registry, version and link editing, version delivery, link following and transclusion following across a distributed network (10, 12); now released under open source at udanax.com), with extensive protocol documentation (29, 30); previously-secret internals are now being published (26); remaining work to be done (31)

• xu92 (formerly called Xanadu 92.1)-- 1988-92 design (9); by Miller, Tribble and Pandya based on K.Eric Drexler's "ent" data system invention (26) and programmed in a small common subset of Smalltalk and C. While full system did not run, the Ent mechanism (ent80) was said to be robustly functional. This code is now released as Udanax Gold under open source at udanax.com

* hcoin96-- HyperCoin(tm) gateway micropayment system for Internet microsale of transpublished content portions (56, 57); implemented as working prototype 1996 by Andrew Pam; patented 2000 (58)

* osmic97-- demonstration of branching referential versioning (48, 47); designed 1996 (49, 50), implemented 1997 by Ken'ichi Unnai in Perl (server) and elisp (client), downloadable (51, 52); implemented again 1999 again under the direction of Yoshihide Chubachi (OSMIC renamed INLUV)

* zz97-- prototype of ZigZag® hyperstructure system implemented by Andrew Pam (from a design completed ca. 1986).

* tpw98-- downloadable and self-installing functioning articulated demonstration of transpointing windows to run under Microsoft Windows; implemented by Ian Heath, based on U.Southampton Microcosm engine

* hts99-- transpublishing server prototype (46) using transclusion formats for ordinary browsers, running from document database; implemented 1999 under direction of Yoshihide Chubachi; design (44); current implementation front page (46)

• xu99-- Floating World design (large preliminary design document 78 (then called ZX)): structures for interactive multidimensional graphics with discrete ZigZag backbone and xu88 fluid-media model

* gzz00-- "Gzigzag", a version of ZigZag® written in Java under the

direction of Tuomas J. Lukka at Jyväskylä University in Finland; available at www.gzigzag.org

* CosmicBook00-- CosmicBook ™ parallel hypertext system implemented by Ian Heath in 2000. Free CosmicBook is reader now deployed from xanadu.com/cosmicbook/

• VLIT01-- a virtual file (content list or EDL) to be opened in Internet browsers, to be fulfilled from rightsholders as required.

Appendix L
1999 Disclosure of Technical Secrets

xutech (d10z
02.01.10

To Xanadu home page
To Ted Nelson home page

A Joint Disclosure by Udanax.com and Project Xanadu
as of August 23, 1999
to accompany our presentation at the O'Reilly Open Source Conference

Xanadu® Technologies-- An Introduction

[Updated to reflect what was actually said at the meeting, with clarifications of the illustrations that were shown. Further citations will be found in our ACM review paper, "Xanalogical Structure: Needed Now More Than Ever" at http://www.sfc.keio.ac.jp/~ted/XUsurvey/xuDation.html. See also our book *Literary Machines*.]

WHAT THIS IS

Udanax.com (formerly XOC, Inc.) have agreed that the software until now referred to as "Xanadu" (particularly Xanadu 88.1 and Xanadu 92.1) is now to be released as Open Source under the X license. All data structures previously held as trade secret by XOC, Inc. are deemed to be no longer secret and all associates and past members of Project Xanadu and XOC, Inc. are released from their non-disclosure agreements with respect to these strucutures.

This page, (and the other accompanying pages at xanadu.com and udanax.com), are the first publication of what have been called the "Xanadu Secrets", in particular, Enfiladics and the Ent.

The shortcomings in this presentation, which are considerable, are entirely due to a lack of time and resources to document what we have done. However, we will endeavor make available all the technical material we can find, either on the Web or through paper publication. Some of these materials are scattered and have been brought back together with considerable effort.

All the structures and algorithms discussed here were disclosed at the Open Source Conference (to the extent that time allowed), and are being made available on the Udanax Website thereafter, in addition to this page at xanadu.com.

The two principal bodies of Xanadu code from 1998-92, until now called Xanadu 88.1 and Xanadu 92.1 respectively, will be packaged for distribution at the Udanax site under the names "Undanax Green" and "Udanax Gold". We desire that development of both shall go forward according to the Open Source model, with contributions from programmers of good will around the world.

BASIC CONCEPT

The principal insight was in place by the end of December 1960. It was this: if text and other media are maintained as referential structures, the resulting structure will have numerous powerful advantages over merely moving the contents around.

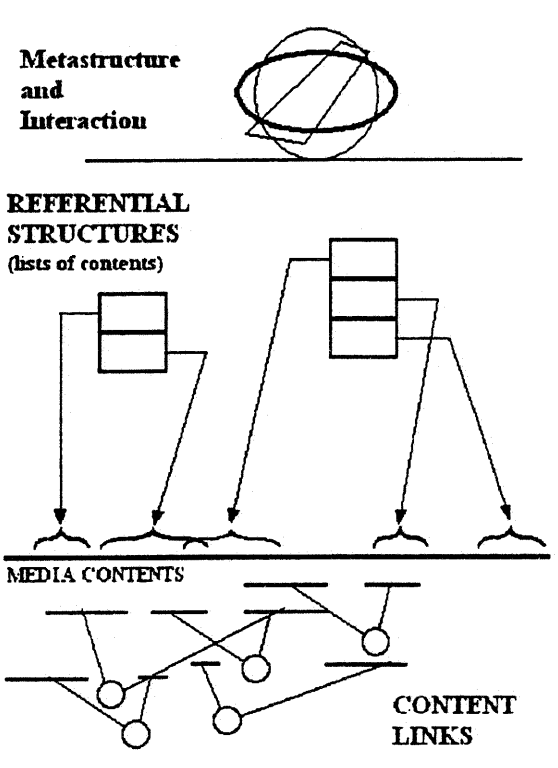

The links (unlike the later links of HyperCard and HTML) are between sections of content.

The advantages of this structure include:

LINKS CAN ADHERE TO CONTENT: if the content moved, the links are still connected.
ANY NUMBER OF OVERLAPPING LINKS MAY BE CONNECTED TO THE SAME MATERIAL.

CONTENT CAN BE MOVED WITHOUT BREAKING LINKS
Therefore we may have a NEW COPYRIGHT MECHANISM:
DELIVERY ON LINE, WITH SALE OF PORTIONS BY
INDIVIDUAL RIGHTSHOLDERS. This could make possible a
new copyright mechanism, with unrestricted re-use without
breaking either links or copyright.
(Moreover, TRANSCLUSIONS CAN BE SHOWN BY
ADDRESS COMPARISON.)

All these ideas are discussed elsewhere, esp. in *Literary Machines*, *The Future of Information* and at Ted Nelson's web pages.

While this idea has always been simple to state, the details must be reconciled carefully, and have resulted in a number of different models. (There are always difficulties in resolving details to a clean design.) A very simplified version was published in 1965 (Nelson, "A File Structure for the Complex, the Changing, and the Indeterminate". Proceedings of ACM Twentieth National Conference, 1965.)

Nelson first presented this idea to a technical group at Brown University in 1968. (That version was called SNP, for "Sexus-Nexus-Plexus", divided into the illustrated levels: content, referential document lists, and links.) The others called the design "raving", so it was dumbed down to one-way links. When asked how the user would get around, Nelson suggested a stack of locations visited. The resulting implementation was the Brown University Hypertext Editing System (CITATION TO CHECK: Carmody et al., "A Hypertext Editing System for the 360", in Faiman and Nievergelt (ed.), *Pertinent Concepts in Computer Graphics*, 1968).

This dumbdown cast a long shadow. While the exact pathways are difficult to delineate, it appears to have established the design of the World Wide Web and the browser as we know it.

VISUALIZATION

In any computer project, the problem is always finding constructs which are sufficient but clean. We published these illustrations in 1972 (simulated in the pictures with cardboard and clear plastic above a Selectric keyboard) to show an extremely simple way of showing two-way links and transclusions.

The idea was of course to see the connections in their two-way reality. The connection-lines should go between the *contents* of the windows, so that either window may scroll or be moved independently.

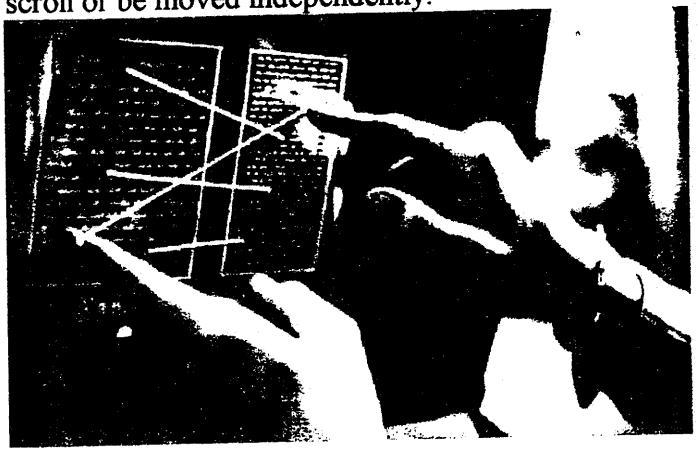

Unfortunately, due to the politics of standardization and other social processes of the computer world, the windowing standard became that later designed at Xerox PARC, which mimicked paper instead, and could not show connections. Furthermore, 1-way links became the prevailing concept; such links could be followed only from origin to target, and never displayed side-by-side.

EARLY TECHNICALITIES

The first eleven years of Project Xanadu was Nelson either by himself or with an occasional single collaborator.

To implement the sort of interconnection shown in the photographs above, various technicalities were devised. This illustration is from a 1971 document, showing the steps for sideways intercomparison (as shown above) in a relatively conventional file representation.

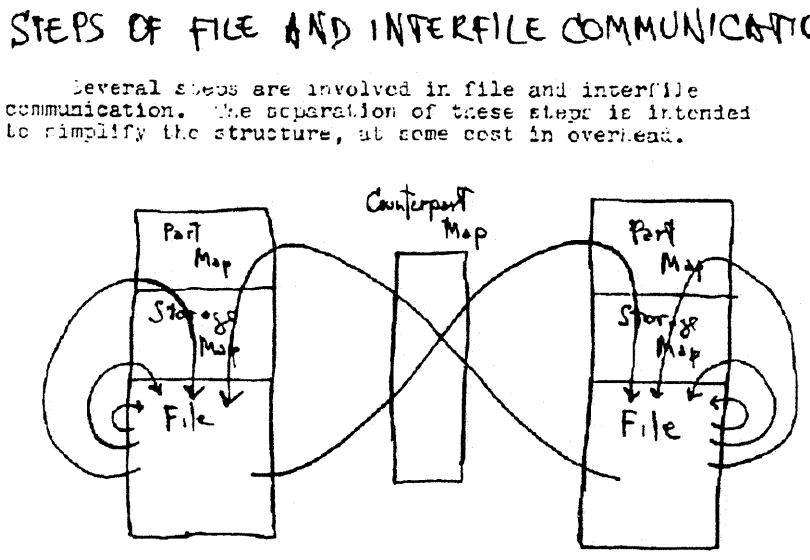

Data structures to suport this were to use parallel files linked in parallel ways corresponding to the parallel visualization:

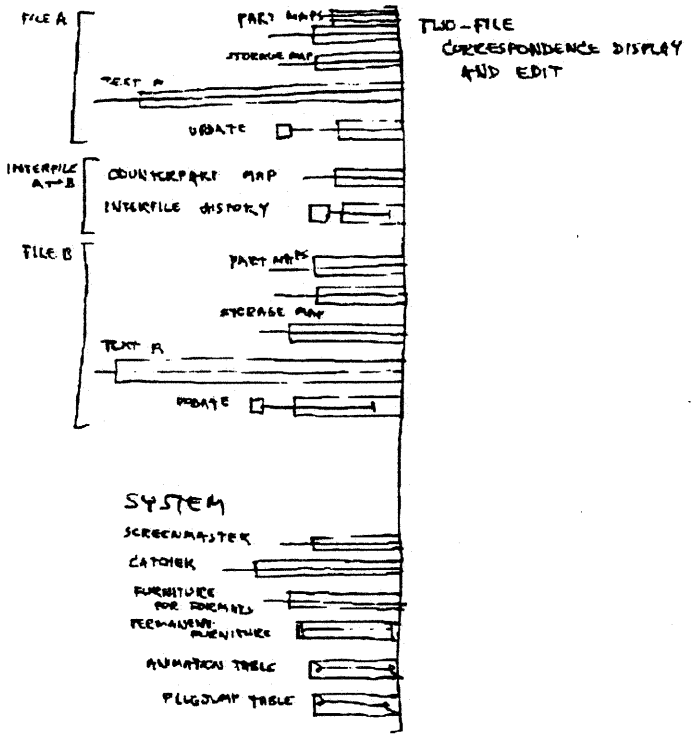

However, it was becoming obvious that more generalized structures would be necessary, to support links and transclusion among documents and versions taken more than two at a time.

A BRIEF HISTORY OF ENFILADE WORK AT PROJECT XANADU

The first enfilade (the "model T") was invented/discovered by Nelson in 1971-2, with the assistance of Jonathan V.E. Ridgway and the late Cal Daniels. (To the best of our knowledge, it was about eight years before others published similar

methods.)

***ENFILADES** [the ***name-with-asterisk** denotes content which was previously trade secret information of XOC, Inc.]

"Enfilades" are a class of data structures (and of course their corresponding algorithms). While a number of enfilade discoveries have since been published by others, the underlying generalities found by the Xanadu group have not, to our knowledge.

As widely suspected, enfilades are principally tree structures, but with a number of additional key features of interest.

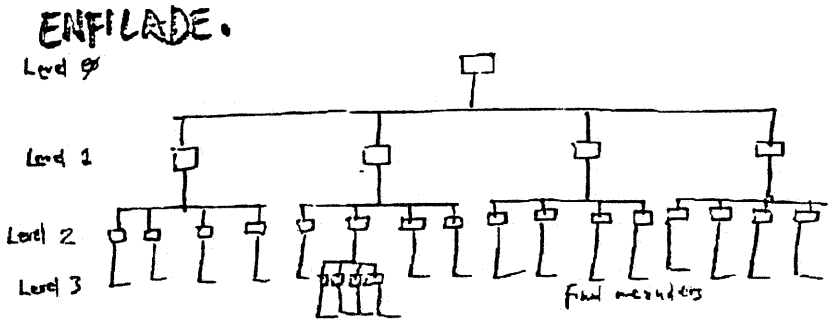

***THE MODEL T ENFILADE**

The Model T was originally designed as part of the JOT project, a subsystem of Xanadu design for editing large bodies of text on cassette tape. (Only accidents prevented the completion and release of JOT in 1973 as the first personal word processor.)

The text itself would not move, but pointers to it would be rearranged.

In order to function with large bodies of text, a tree structure for pointer-based editing was devised. An intrinsic design intention was to allow the tree to grow to any size by systematically maintaining a subrepresentation of the tree in core--in other words, system-supported caching particular to this method.

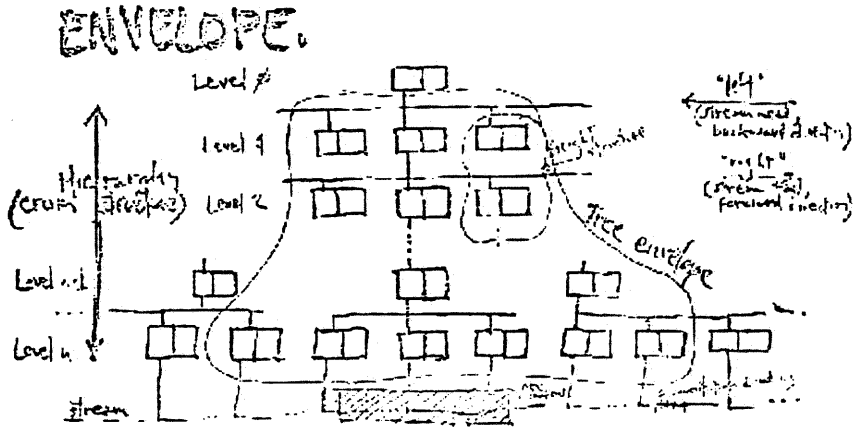

The central feature of technical interest was that each pointer contained a count of the numbers of characters below (a field we called the WID, since it showed the width of contents referenced underneath). This is indicated in the following picture as the second square of each pointer.

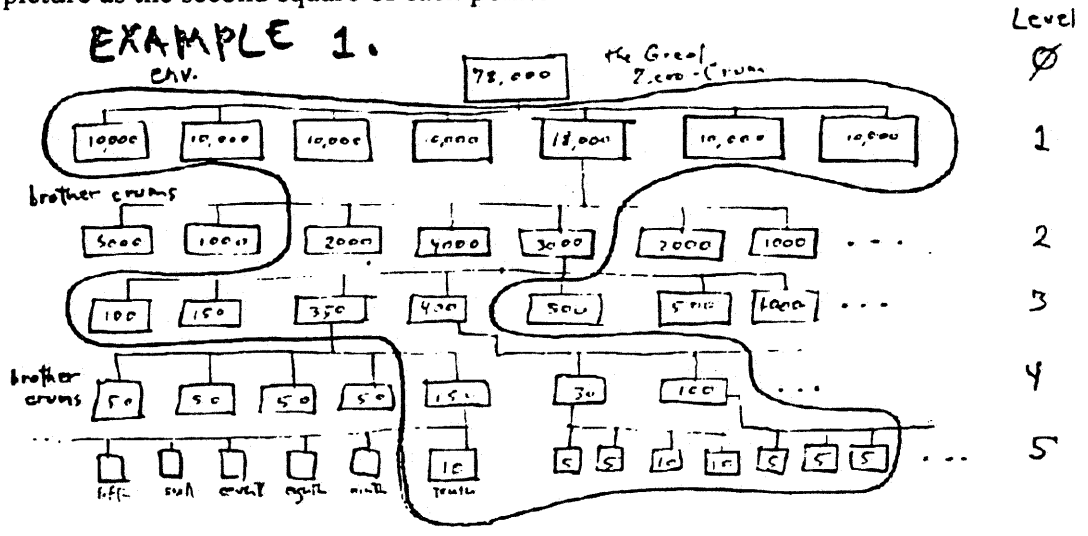

WIDs in amoeba of the Model T enfilade

The lowest-level pointer (level n in the above illustration) pointed directly to a text segment. Its WID was the length of that segment. Above level n, the WID field summed the WID fields below.

This sum of WIDs propagated up the tree to the root pointer, whose WID therefore contained the text length of the entire document. This arrangement made it possible to go directly to any location in the document, by evaluation of the WIDs on the way down to the final content.

EXAMPLE OF USE

Today's word processors typically offer only rudimentary rearrange, by deleting text into a buffer (the misleadingly so-called "clipboard", rather than the richer forms of rearrangement in constant use by prose authors). We designed the 1972 system for rearrangement of two sections to be exchanged, either consecutively or around an unmoving section. (These were respectively called the Switcheroo, with three cuts defining two sections, and the Switcheroonie, with four cuts defining three sections.) An author was to delineate the cut-points with exclamation points (a convention later adopted by Word Perfect). These cuts would then be transferred to the structure of pointers and be applied to the contents below.

The following example shows how three cuts would determine a rearrangement.

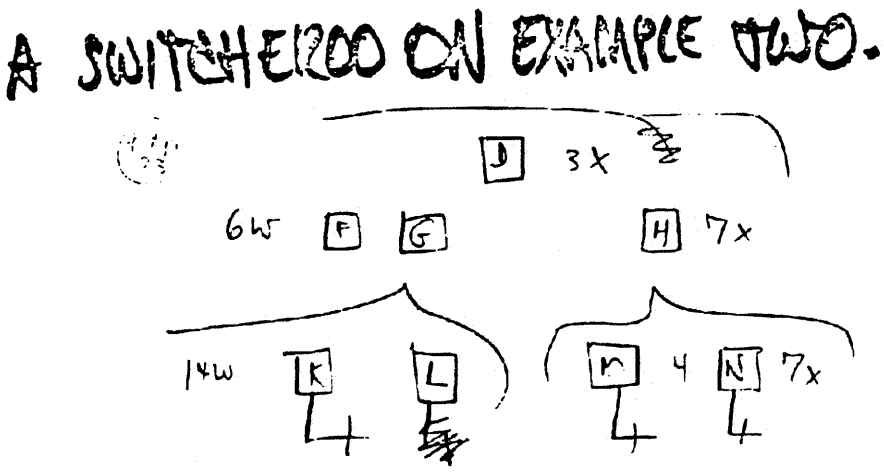

Internals of rearrangement, Model T enfilade, 1972.

Letters here identify reference pointers ("crums") being kept in core.

THE CRUM TABLE

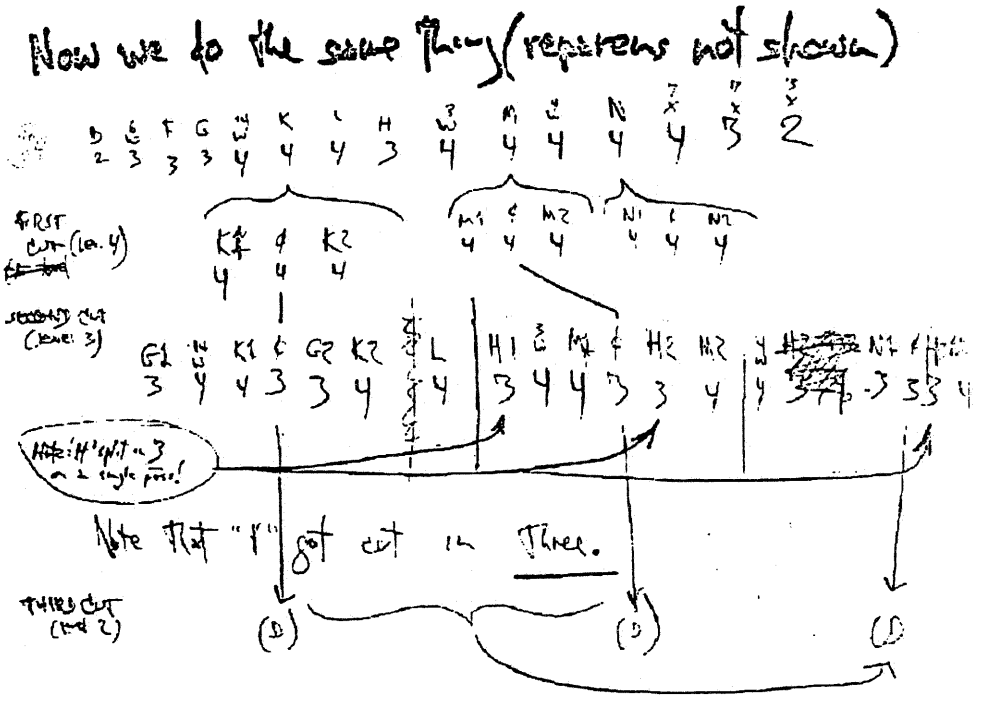

Crum Table enacting rearrangement in the previous illustration.

The crum table shows how the rearrangement is enacted within a buffer containing and identifying the crums. A numeral is associated with each crum, showing its level in the tree structure of the enfilade, thus modelling a subrepresentation of the whole structure. The lowest-level crums (level 4) point to string data. WIDs are not shown.

In this map, the three-cut rearrangement of the previous illustration is mapped into its constituent operations on the crum pointers. The "¢" symbol indicates a cut to be made in the crum, specfically to crums K, H, M and N, which we see propagated up to the highest necessary level. A rearrangement of the crums (and sending the crum table to mass storage) completes the rearrangement of the text below.

The first two rows of the illustration identify the crums of a subrepresentation of the enfilade, starting with D, a crum at level 2. Passing rightward, six crums are omitted from the map; then F, a level 3 crum (not involved) and G, a level 3 crum, where the work begins. Below G are 14 level-4 crums (omitted), then K, a level-4 crum involved in the operation, and its succesor L, not involved. We see H, a third-level crum involved in the rearrangement, above M; four other crums are not shown, N (involved in the operation), and then other crum-counts of crums which are omitted.

In the third level of the diagram, showing only a part of the table, we see the cuts propagate to the bottom level as the rearrangement is consummated. Sections of crums are exchanged, after which the "¢" symbols may be removed from the crum table.

To enact the operation, K, M and N are cut. In the method illustrated here, the "¢" sign is used as an actual marker in the crum table. Crum K is replaced by crum K1, the "¢" sign and crum K2. Crum M is replaced by Crum M1, the "¢" sign and crum M2; crum N is replaced crum N1, the "¢" sign and crum N2.

In the last level of the diagram,

Depending on the depth of the tree and whether the upper crums contain all the affected contents, the cuts may have to propagate to additional levels, which can be easily done by iterating the same method.

*THE BARUS CONJECTURE

William Barus, then a philosophy graduate student at the University of Chicago, generalized this basic enfilade in surprising directions. He conjectured that version management could be handled by separately maintaining two sets of pointers to the content-- the original input sequence, and a current version. This approach led directly to the other conceptual breakthroughs.

Based on Barus' discoveries, the Xanadu group gathered for the "Swarthmore summer" of 1978. This led to
1) the discovery and elucidation of General Enfilade Theory; 2) the design of the Xanadu 88.1 system (Udanax Green); 3) the discovery of the Ent; and 4) Xanadu 92.1 (now Udanax Gold)

*GENERAL ENFILADE THEORY

(The reader should understand that this brief introduction is not the whole of the theory, and is referred to the originators or to Roger Gregory.)

General Enfilade Theory was discovered and elucidated within the group during this period (though still not published) by Mark Miller and Stuart Greene (now Stuart Grace). They discovered that enfilades had two basic properties which could be independently tailored to create powerful data structures as needed.

The two properties were called WIDativity (upwardly-propagating properties, such as the WID field of the Model T enfilade) and DSPativity, which was also present (though unsuspected) in the Model T enfilade. DSP, standing for "displacement", applies to properties which propagate cumulatively downward, being imposed from above.

Example: in computer graphics, a WIDative property (upwardly cumulative) is the bounding box, which grows as the tree included becomes larger. A DSPative property is matrix rotation, such as that of a finger of a hand on an arm, where the movement of a single bone of a single digit is a matrix product of all the motions farther up the tree.

What was the DSPative property in the Model T? It was the sequencing of the pointers, which sequenced the text below in a tree of downward imposition.

Both WIDative and DSPative properties must be associative, i.e. for such a property p

$$A\ p\ (B\ p\ C) = (A\ p\ B)\ p\ C$$

because there is no telling how the next version will be edited. However, WIDative properties must be associative horizontally (like the sum of the WIDs in the Model T enfilade) and DSPative properties must be associative vertically (like the positions of the crums in the Model T enfilade).

General Enfilade Theory permits the creation of custom enfilades by the suitable selection and design of WIDative and DSPative properties.

XANADU 88.1 (now "Udanax Green")

The Xanadu 88.1 system was designed by Roger Gregory and Mark Miller in 1981, and essentially completed in 1988. This system, described in our defining book *Literary Machines*, searches a vast address space such as the Internet, creating a distributed representation of the system on many servers. Each server manages its own dynamically explored subrepresentation of the whole. Each node may subcache portions of the docuverse as required. The whole space may be searched with unusual efficiency for overlapping content links and transclusions.

Consider the linear address of the docuverse, or a subset, as a line. (See illustrations in *Literary Machines*). This address space is measured in tumblers, which are multipart numbers of the form 0.zzz.yyy.xxx ..., where the fields "zzz", etc., are any integer. The zero is convenient for the calculation. The first field zzz represent a node, the second field yyy represents an account (perhaps author or company), the fourth field xxx represents a particular work, the fourth field wwww represents an address within the work, and the fifth field vvv represents the element type. What is particularly interesting is that stretches of the docuverse are represented by a starting tumbler and a difference tumbler, manipuated with unusual numerical address operations based on transfinite arithmetic. (This is presented in *Literary Machines*.

***PREVIOUSLY SECRET INTERNALS OF 88.1/U.Green**

***The Istream.** Consider the docuverse order of content pieces (sometimes the local order of arrival and storage) as the invariant stream, Istream or I. (It is not going anywhere, and will be rearranged in place.)

***The Vstream**. A current permutation of a document is represented by another ordering of the same contents. This is called the variant stream, Vstream, or V. This rearrangement V is represented as a permutation matrix.

It was explained in *Literary Machines* that tumbler arithmetic is used for calculating spans of addresses. Actually tumbler arithmetic is used for address matrix operations throughout the system.

***ENFILADES OF U.GREEN**. There are three interacting enfilade structures: the Granfilade, the Poomfilade, and the Spanfilade.

***The Granfilade** (Grand Enfilade) manages the content of the whole address space, similar to the function of the Model T enfilade. As new content is added throughtout the docuverse, its presence percolates upward as counts in the whole system. (How this is distributed is another matter.)

***The POOMfilade** (Permutation Of Order Matrix Enfilade) is local to a document. It provides the I-to-V and V-to-I transform-- that is, given any location of contents in the address space (an I-address) it finds the corresponding current address of that element (V-address), and vice versa.

***The Spanfilade** (Spanning Enfilade) is where the real work of the system occurs. It provides V-to-V transforms, making it possible to go from a Vstream address in one document to the Vstream address of the same content in another document. This is the address comparison for both links and transclusion spoken of earlier. The spanfilade is also intended to provide the sieving function that filters out the irrelevant connections.

The Udanax Green code is in C.

Udanax Green ran on a small scale, but development was stopped in favor of an improved design, the less-finished Udanax Gold, based on Drexler's Ent.

*THE ENT

The Ent was invented/discovered/designed by K. Eric Drexler (discoverer/inventor of nanotechnology and head of the Foresight Institute). The Ent, named after the walking trees with great memory in J.R.R. Tolkien's *Lord of the Rings*, is a data structure and algorithm plex which manages complex versioning of arbitrary objects with low overhead and great parsimony, representing each version simply by a handle into complex of data.

XANADU 92.1

Xanadu 92.1 (now Udanax Gold) was architected by Mark Miller (who later acknowledged as his co-architects Dean Tribble and Ravi Pandya.).

***PREVIOUSLY SECRET INTERNALS OF 92.1/U.GOLD**

This consisted of the Ent and various object-oriented wrappers for it, seeking the most robust, efficient and generalized version. Xanadu 92.1 was still unfinished when Autodesk dropped the Xanadu project in 1992.

Xanadu 92.1 system was based on wrapping the Ent, mapping it to a highly-generalized multidimensional coordinate space for arbitrary data, and encapsulating requests and deliveries for efficiency.

Design of 92.1 began in 1988. At that time, object-oriented programming tools for large projects were in poor shape. The group chose to maintain identical parallel code in a carefully-chosen mutual subset of Parc Place Smalltalk and C++, in order to obtain certain benefits from each.

The U.Gold code has satisfied certain tests but is not considered to be usable as yet. However, the kernel Ent routines, wrapped by the rest of the system, have worked well since about 1990. It may be an interesting project for someone to endeavor to extract the Ent code from U.Gold (the name for the distributable version of Xanadu 92.1), in the same manner that the "curses" library was extracted from Bill Joy's vi text editor.

AVAILABILITY

The code will be available from the udanax.com website, with a release date of 23 August 1999.

LICENSES AND TRADEMARKS

Release will be under the Berkeley II open source license, requiring no

acknowledgment of use.

To avoid confusion with present work at Project Xanadu, the names "Udanax Green" and "Udanax Gold" will be given to these packages and their ongoing versions. However, this software was called "Xanadu Code" in the past, and of course any reference to it in the past tense may correctly refer to it under the trademark "Xanadu". It is just not being called Xanadu from now on, as it spins off into new adventures.

Appendix M
Xanadu Designers

To Ted Nelson Home Page

Appendix M.
Xanadu Designers.

This is slightly extended from "**Xanalogical Structure, Needed Now More than Ever: Parallel Documents, Deep Links to Content, Deep Versioning and Deep Re-Use**" at http://www.sfc.keio.ac.jp/~ted/XUsurvey/xuDation.html

Theodor Holm Nelson, Project Xanadu* and Keio University

Principal designers involved in various versions of the Xanadu designs have included (in alphabetical order)--

Paul Baclace
William Barus
Les Carr
Yoshihide Chubachi
Cal Daniels
K. Eric Drexler
Stuart Greene
Roger Gregory
Edward Harter
Ian Heath
Chris Hibbert
Eric Hill
Hugh Hoover
Rob Jellinghaus
Roland King
Tuomas J. Lukka
Marlene Mallicoat
Michael McClary
Mark S. Miller
Theodor Holm Nelson
Kiyoki Ookubo
Andrew Pam
Ravi Pandya
Bill Richard

Jonathan V.E. Ridgway
Jonathan Shapiro
Marc Stiegler
Johan Strandberg
Dean Tribble
Ken'ichi Unnai
Steve Witham.