What is the nature of language space? John Cage took apart the dimensions of sound space: duration, pitch, loudness, timbre, envelope—what are the dimensions of language space? How do we delineate the geometry of language space if not dimensions?

How do we allow for the totality of possibilities for language space?

What does it mean to bring multiplicities inside the sentence? Inside the thought?

What is the nature of the interactive word object? How should interactive word objects work? How do they interact with one another — both reinforcing and inhibiting?

What does it mean to have an ecosystem of word objects? What does the substrate look like?

What are the possibilities of a form of writing in which any element can be connected to any other? Including feedback loops?

What does it mean to "write native" in interactivity?

E-Poetry 2007, Paris

CONTENTS

PREFACE

The essays and papers in this book were only begun after what has become a lifelong quest was well underway. My earliest experiments in non-linear writing—which I came to call word-nets—were begun in the summer of 1966. These may best be described as polylinear: interlocking linear phrases arcing across a large graphical space (mostly done on large pieces of watercolor paper). This form is not what I had originally thought to make. The original idea, which came to me in a flash, was to make a grid of words that could be read along any pathway of adjacency. (It was only several years later that I discovered that this exact form had been invented no less than a millennium and a half ago in China, as the celebrated Xuanji Tu, also known as The Revolving Chart, of Su Hui—more about which below. Over the years I have come to use Sui Hui's amazing work as a kind of stick to hit myself over the head with at moments of being in danger of thinking myself entirely more clever than is warranted.) Having had this flash of an idea, my first thought was that it would not be until I reached the venerable age of 30 that I would have learned enough to be able to actually do this. But, later that summer I thought: "I'm going to try!" And out came a word-net instead. A couple of years later came my first experiments with diagram poems. Little did I know then that the diagrams would come to be the core of my life's work.

The impetus for the diagram poems came from a term I was hearing from undergraduate friends I was hanging out with in the music crowd: they were excited about something called Tone Clusters. The minute I heard this it was like an explosion going off in my head. I knew immediately: I have to make word clusters! But, right away there was a problem. What part of speech is a word cluster? This question completely brought me up short. Syntax is partly brought by context, but a large part is brought by the word itself. If words are brought together into a cluster—to occur "all at once"—then the parts of syntax brought by individual words will interfere with one another and "not be" a part of speech. What to do. Almost immediately the thought came to mind to externalize syntax, to pull it out into its own channel, to just draw it. So began, in 1968, the Diagram Poems. My recollection is that the first time any of my diagram poems appeared in public

was an installation done at The Kitchen in New York in 1975. Diagrams Series 3 appeared in 1979, and Diagrams Series 4 in 1984. All throughout this time I was doing no theory at all—no essays, no papers, no poetics statement—absolutely nothing that anyone would recognize as theory. (In spite of several requests for a poetics statement on the Diagram Poems.) At the time I could not have articulated why I felt I simply couldn't do poetics; I simply had a sense that I had to be at a certain place first, and knew I wasn't there. Midway through Intergrams, I realized I was there. The result was "Openings, The Connection Direct."

I cannot express my attitude toward poetics any more clearly than I did in an interview in this book:

There is an analogy from biology that I find useful as a metaphor for the artistic process that may help explain how I feel about poetics. We all are so used to navigating by senses that operate via direct lines that we can't even imagine what it would be like to function in a world where the primary sense is chemical. Some organisms that orient to chemical gradients use a mechanism called klinokinesis. (I hope I'm getting this right ...) These animals are constantly turning, and the rate of turning is adjusted based on whether the turn took them to more favorable or less favorable conditions. It works well for orienting to chemical gradients, and even though these animals can't "see" the source of their food, by the laws of probability the turning is controlled in such a way that they end up at the points of maximum concentration of the chemical to which they're orienting.

This says a lot to me about the artistic process. One is constantly turning, and the important thing is that fine-grain sense of feedback: this worked, that didn't, this feels right, that doesn't. One can't always directly see hulking above the landscape the exact artistic endpoint, but ultimately the feedback from turnings enables that journey that somehow gets you to the point of maximum energy. Poetics, on the other hand, runs the risk of erecting a line-of-sight structure that sticks out above the landscape. There's nothing wrong with that, of course, but woe unto you if in the process you manage to blast those sensory nerve hairs that make the klinokinesis work.

... having gone a journey by klinokinesis, one may be able to see what the topography was and explain it in a way that is useful to others; I suppose I tend to think of poetics as being a more retrospective activity than a prospective one.

During this time I confess I had a rather bad attitude about hypertext; like many people still today I associated hypertext with nothing more than nodes and links, and felt that didn't really have much to do with what I was doing. In the Spring of 1994 or thereabouts, Mark Bernstein at Eastgate was prevailed upon to pull together (on short notice) a literary panel for ECHT '94, the European Conference on Hypermedia Technology. He put out a call for papers. I sent a rather sketchy essay-like draft that was more like an extended abstract than a paper—at the time I believed "I don't write papers". To my surprise he replied that a "real paper" would actually be rather nice. So, starting from zero, I ordered the entire set of ACM hypertext conference proceedings, and set to work. Having been told by several people that I should take a look at Cathy Marshall's work, I started with her Aquanet paper, and the result was electrifying. Here was a hypertext system based on relations, not nodes and links. My diagram notation was inherently relational. Her graphical representation of relations was so close to my diagram notation that I practically jumped out of my chair. If the hypertext community counted Cathy Marshall's work as hypertext, WELL, I could count myself in on hypertext after all. Out of this came "Navigating Nowhere / Hypertext Infrawhere"—and all the papers that followed.

So we come to the present moment, which in a way is bittersweet. In use by billions on the World Wide Web, we could say that hypertext has conquered the world. But has it. Most web pages have no more than a scattering of links—to whole documents. The kind of intense, word-level, inside-the-sentence connectivity pioneered more than 1500 years ago by Su Hui is almost nowhere to be found. The excitement of the founders, those like Vannevar Bush, Doug Engelbart, Ted Nelson, that we were engaged in a great adventure in creating a medium of thought itself, seems to have dissipated completely. That magnificent functioning wonderful anarchy called the Internet is being corporatized before our eyes—as those from the Western United States would say, the range is being fenced in.

Much as I might admire the incredible achievement of Su Hui a millennium and a half ago, I have to admit, with tears in my eyes, that hers was a moment lost. In Su Hui's day, the world was not ready for hypertext. So, will we lose this moment also? I fear.

I can't resist closing with a personal story about a moment of transformation in my life, one of those moments where you are changed forever. It came when as a very young poet I was reading Charles Olson—whom I never met and never even heard read live. As I recall, the book was The Mayan Letters. I came to a passage that affected me so violently I felt as if I had been struck in the breastbone by a sledgehammer, so hard I literally had the feeling of flying through the air. The passage was very simple. When I tell it to you you'll be puzzled that I had such a violent reaction. It said, simply: poets should study cosmology. Far be it from me to be interpreting Charles Olson, but I can certainly speak to how I took this phrase. I read it to mean: poets should study cosmology. Cosmology is the study of the universe. All of it. Meaning: there is no boundary to what, as a poet, you may be called upon to learn about. The reason I had such a violent reaction to this phrase is that I realized I was a pathetically narrow person, and that had to change. It is a problem I struggle with still.

So, here I write, a millennium and a half since Su Hui put words together more intricately than anybody since, and after a handful of decades of work I can look at the mountain, hypertext as a medium of thought, and think: I've gotten to base camp. The mountain is still up there.

Cosmology: pass it on. Let's go.

Acknowledgments

The work in these pages would not exist without the work of many wonderful people. To my forebearers: Charles Olson, Robert Creeley, Robert Duncan, Denise Levertov, Jackson Mac Low, and many more; to Judy Malloy, Michael Joyce, Stuart Moulthrop, John Cayley, and Philippe Bootz, along from the start on this adventure; to Cathy Marshall, the mother of spatial hypertext, who inspires continually and gave me the place to shoot from; to Mark Bernstein, who

endured various tantrums and taught me that if I didn't like the tools I had, I had darn well better learn to talk to the people who could fix that; to Sandy Baldwin, who eggs me on to get to the heart of the matter; to Loss Pequeño Glazier, who first brought us E-poets together; to Chris Funkhouser, for more patient attention than I deserve; and most of all to John Cage, infrawhere, the most magnificent Opener who ever lived; and to MJK, the co-, completely.

Grindstone, Pennsylvania
February, 2013

OPENINGS: THE CONNECTION DIRECT

PERSONAL NOTES ON POETICS

OPENNESS

Poetry is not a circumstance of language. Rather, any possible circumstance of language is a possible circumstance of poetry. It is the job of the poet to invest that circumstance with energy. It is the job of the receiver to be open-minded about what circumstances of language may constitute poetry. This is the exact analogue of the idea that the domain of music is anything which may be heard, or that the domain of the visual arts is anything which may be seen. The page may be a wall or a computer screen or a street or a floor with words glued together in a pile so that not all of them can be read. This is not meant in any way to disparage the traditional page. If there can be such a thing as a conscientious avant garde, then surely its purpose must be to expand the field of possibilities for making art, not to replace the existing set of possibilities with a new one, equally narrow. The house of poetry has room for everyone.

ENERGY TRANSACTIONS

It is a common stereotype that the arts "are" communication processes. Communication is a wonderful thing; no one should put it down. Like love, it doesn't always happen when you want it to. What is pernicious in the arts is the view that when communication has failed, all has failed, that there is nothing but communication. In fact there is a layer that underlies communication: the energy transaction layer. The artist is presumably a person who is able to take energies and make them available in concentrated form. The receiver is presumably someone with certain energy needs. What is important is that the transaction take place: that the energy is transferred. (What the energy "is", where it comes from, how it works—these are all questions on which we will all, of course, differ.) Where

the artist has an exactly clear view of just how the energy transaction should take place, and it does in fact take place that way, then we call this communication.

But imagine you have just walked into an art gallery. All about the space are works that you don't understand at all. You find them irritating, perhaps infuriating. You "get nothing" from them. You walk out of the gallery in complete disgust. The artist has not communicated with you at all—you are certain of this. Then a strange thing begins to happen. You begin to notice that somehow, you are simply seeing much more sharply than you usually do. You find yourself at a heightened state of attention. You find your thoughts making connections that hadn't happened before. Your life seems suddenly more clear. As opposed to that dreadful mess in the gallery!

This is the energy transaction at work—in a way that perhaps neither artist nor viewer "intends" or understands. And yet, by being energized, by being brought to a heightened state of attention, something useful has happened in your life. Something useful to you. Not an injection or gift of someone else's wisdom, but a connection that was there for you to make all along, something entirely yours, a connection that sprang forward with the impetus of the energy of the works in the gallery. Let us admit that communication failed here, but the energy transaction worked.

The artist is one person, but there are many potential receivers. Simple arithmetic announces who is important here: the receivers! For an artist interested in energy transactions, the purpose of art is to jog the receiver so that resources already there in the receiver's mind are brought together in a productive way. You may not like this idea. You may find it like asking a question and being greeted by another question in response. You may wish for an injection from somewhere outside of yourself. If you do, you may not be happy with an energy that operates on your own resources without giving you new ones.

Communication, by definition, means being specific about what energy transactions can take place. But an art that focuses on the energy transaction layer itself as the primary layer should seek to maximize the energy transactions that can take place. This means the artist should not stand in the way of her/

his own energy transactions. For an artist who is not specific about what energy transactions should take place, there is no "thing" to be communicated.

The distance away from the energy transaction layer at which an artist wishes to place the focus is an artistic variable; like any other artistic variable the artist may choose to vary it over a life's work or within a single work or not at all; the artist may seek to make it clear or ambiguous, or even "flip" as our perceptions can flip when viewing an optical illusion.

Sometimes the energy transactions do not come off, but a third party can help to bring them about. This is the proper role of criticism. This is the only proper role of criticism. Critics who actively seek to prevent energy transactions from taking place because they consider them of no value are harm-doers performing destructive acts, and should be labelled as such, like other vandals.

NON-POSSESSIVENESS

The communication stereotype goes hand in hand with the compulsion to possess "the thing communicated." Possessiveness carried to extremes can have unpleasant side effects. In order to best receive works intended to operate primarily in the energy transaction layer, it may be necessary for the receiver to make the effort to be purged of possessiveness. Some potential receivers will not want to do this. They will find the work difficult. The work may indeed be difficult, but nowhere near as difficult as giving up possessiveness.

Imagine you are in a primeval rain forest, surrounded by sights and sounds that are completely unfamiliar. All at once the most amazing bird you have ever seen flies by. Its iridescent colors look like nothing you've ever seen. Alas, you see the bird only for an instant, for a flash so brief you can hardly be sure you saw it. Then it's gone. Perhaps you will feel that you simply must get a full unobstructed view of that bird. You go crashing through the forest trying to find it. You become manic about it.

Or perhaps you become very still, surrendering to whatever the forest chooses to show you. You would be thrilled to see the bird again. But you know you may not. You are energized by that one brief moment of having seen it, as you try to be energized by every moment the forest has to offer. You move through the forest slowly, becoming part of it. You see all of it, even though of course you see only a tiny part of it.

Which strategy is more likely to get you another sighting of that bird?

NON-LINEARITY

Alas, McLuhan got it exactly backwards: speech is an inherently linear medium, writing is a medium with an inherently non-linear potential. We think of writing as linear only because writing is such a young thing: we haven't figured out yet how to tap even a fraction of its power; mostly we use it as simply an alternative medium for nearly the same languages as are used for speech. In short, the supposed linearity of writing comes simply from a lack of inventiveness in using writing only for the same kind of language that might be spoken. So we must begin by understanding the linearity of speech. (Or Sign, for that matter; the issues for both are identical.)

Speech is an activity in which speaker and listener are constrained by the requirement that comprehension must take place in real-time. (But word-time is not quite real-time—~~it~~ has an oddly retroactive character. We "hear" what the mind has retroactively decoded from the previous so many fractions of a second as though we were hearing it exactly as it happened.) It is the inexorable linearity of time which makes speech linear. (Or makes performance of any kind linear, for that matter.)

A simple model for language would have the listener decode the message by going through a series of states; the rules for language would tell, based on what was heard and the current state, what the next state would be. This model has the virtue that it requires very little storage: the listener must only remember the current state. This kind of model is known as a finite state machine. It is

known that there are serious limits to the complexity of structure that can be built with finite state machines. (In particular, a finite state machine cannot handle "self-embedding" structures. A sentence that has another sentence in the middle of it, e.g. "His statement that he has nothing to hide will not wash, and I told him so," is a self-embedding structure.)

Only fairly simple languages can be handled by the model of a finite state machine. When more complexity is required, we will need a model with a much more exacting requirement for storage. Sentences in natural languages require the listener to fit together pieces of what is heard that may be widely separated in time. The method by which cues are embedded in speech as to how the storage should operate is called syntax. You could say that syntax is "speech's way out" of the linearity imposed by time.

But writing is vastly different. First, there is no constraint imposed by time. The reader may reread—~~or~~ skip around—~~as~~ many times as needed to feel comfortable. Second, a written document, unlike a spoken performance, contains its own storage. The storage burden does not fall so completely on the reader as it falls on the listener. In writing, space replaces time as the fundamental dimension set for text as opposed to speech. Complex links between parts of a written text separated in space may simply be drawn directly. The method of directly, graphically linking the pieces of text connected by a relationship can be used for syntax itself: Direct Access Communication—~~as~~ opposed to speech, which may be called Synchronous Sequential Access Communication.

And yet we seldom find works written to directly exploit these capabilities. Instead, writing tends to be used merely to freeze-dry speech. No wonder there is such a strong feeling in the poetry community that the spoken word is the primary medium, has far more power than writing, that to know what a poet is "really up to" you have to hear the poet read. We haven't yet learned to start writing.

HYPERTEXT

The term "hypertext" was coined by Ted Nelson more than two decades ago to describe a way of organizing text aided by a computer that allows the reader to follow links as well as simply "read forward." "Traditional hypertext" allows a non-linear organization to be superimposed on an otherwise linear document. Or pushed one step further it allows a non-linear organization to be imposed on a locally linear substrate.

Hypertext does not go nearly far enough. The non-linearity should be extended all the way down into the fine structure of language. Syntax itself can operate through the same kinds of operations as the hypertext link.

INTERACTIVITY

Hypertext is interactive in that the user makes choices concerning which links to follow. In some hypertexts the user may additionally create links as desired; other hypertexts are "read-only": the user may follow links but not create them. For read-only hypertexts, interactivity is mainly concerned with navigating the link space of the hypertext. There are numerous unresolved questions in hypertext that are the subject of lively investigation. How best should links be created? How can the structure of the link space be presented to the user without the user getting lost?

Hypertext-like navigation can be used with direct access communication to achieve a virtual page of arbitrary size and complexity.

Navigation is of course only the simplest thing that may be done with interactivity. Interactivity may be used to allow the user to change the entire structure of the text. Interactive text can be said to behave; the only limits to the possible degrees of complexity of this behavior are the general limits to the complexity of behavior of computer programs. (See below, "the animate object".)

JUXTAPOSITION (= STRUCTURAL ZERO)

The use of juxtaposition—superimposing elements with no structural relationship whatever—is taken for granted as a possible artistic device in music and the visual arts. In poetry it presents a very profound problem. Juxtaposition is structural zero. To paraphrase Cage's old critique of the twelve-tone system, syntax is a vocabulary of structural descriptions which has no zero. In a traditional sentence, every word has a structural role, every word has a structural relationship to every other word by virtue of where each word is in the syntactical structure. It is not possible to have a sentence in traditional languages where words are just "there together" with no structural relationship at all. So the poet wrestles with a difficult dilemma: forego juxtaposition, or forego syntax.

In direct access communication, the burden of syntax is removed from the words and carried by the medium itself, through direct links, be they graphical on paper or hypertext links. This allows the use of juxtaposition and the kinds of structuring provided by syntax. The syntax of direct access communication is a syntax that allows for zero, that allows for elements that are juxtaposed without structure to be combined into a larger structural whole. The dilemma over the use of juxtaposition is solved. Syntax becomes an option but not an obstacle.

SYNTAX WITH ALL SLOTS OPEN

A visual, diagrammatic syntax is a syntax with all slots open. Any point can be connected to any other point just by drawing the link. Poetry is given the openness that has been taken for granted in the other arts for decades, without giving up the richness that syntax provides as a vocabulary of structural descriptions.

The ability to draw syntactic links directly makes it easy for direct access communication to create syntactical possibilities that are difficult or impossible for traditional grammars. Example: the feedback loop. Feedback loops are among the most ubiquitous and fundamental structures in nature. They are also notoriously nasty for theories to deal with. Traditional grammars do not allow for feedback loops. There may be a loop in the sense that a grammatical rule is revisited,

but in mapping out the syntactical parts and their relationship to one another an actual element of a sentence is not structurally revisited. But the openness of a diagrammatic syntax makes this easy. A link indicating a predicate may end up back at an element that's part of the complex being predicated. The eye can see there is a loop, can take in the whole loop as a structure. Things lead back: we all know sometimes life works this way.

A grammar permitting feedback loops would be impossible for a computer to deal with. The computer would hang in the loop, would not realize there is a loop, would be able to form no gestalt for the loop as a whole. Even the mind would have trouble with a feedback loop in the medium of speech: the linearity of time makes it too hard to go back, and back again, and form the gestalt of a loop.

Another structure made easy in direct access communication is the internal link. An internal link is a link between an element and a larger complex in which that element participates. (This is a form of loop, actually.) Consider a clause, and the relationship between the noun and the whole clause. That relationship itself—the role the noun plays in the clause—is not available in traditional syntax as a syntactical element. But in a visual syntax, an internal link is as easy to draw as an external link.

VOCABULARIES

The historical "first problem" for computer poetry has been how to get the words into the computer. Of course this won't be a problem much longer: computers will all come with good dictionaries, the words will already be there. But in the past, to use a computer with poetry has meant first getting the words into the machine. But putting all the words of a natural language into the machine is a huge undertaking. So the computer poet had to first give the computer a vocabulary, a vocabulary more restricted than the whole of the language, more restricted than the totality of words the poet knows. This has been a major stumbling block for many poets who might otherwise have worked with computers. The nasty word here is "restricted." Poets do not like feeling restricted. "Vocabulary" is not a conception many poets find congenial with their poetics. (Jackson Mac Low

comes to mind as an example of someone who, to the contrary, has worked with vocabularies as an element of his poetics for decades, computer or no computer.)

Is a vocabulary a "closed form" (in the Olson sense)? At least one vocabulary everyone knows is (trivially) not a closed form: the entire language. If the entire language is a closed form then the term "closed form" makes no distinction so we should quit talking about it: everything would be a closed form. So the question must be rephrased: how small must a vocabulary be before it becomes a closed form? We will differ on the answer to this, of course. Personal view: a vocabulary can be amazingly small and still be an open form. That is the hard part, of course: composing a vocabulary which is small but still open.

We should not be bashful about small numbers. (Robert Duncan always said he couldn't count beyond five.) A vocabulary "composed small" will induce repetition in the works composed from it in a way which is musical but not overt.

PRECOMPOSITION

Precomposition—~~c~~omposing, prior to creating the visible/audible/readable "ultimate elements" of a work, a "layer" which affects the entirety of the final work—~~is~~ a venerable concept. Visual artists have been doing it for centuries. The woods had to be scoured for materials to be ground into pigments and a canvas had to be stretched and material prepared for gesso and primer coats applied and then undercoats applied—~~a~~ll before a single square millimeter of the final surface was "painted." Composers, particularly electronic composers, have practiced precomposition extensively. But for some reason, the concept of precomposition seems to be in poor favor among poets. It's as if we are still struggling with a ghost of romantic idealism about the act of composing poetry that looks on precomposition as dirty, somehow. Allen Ginsberg has written explicitly about the act of composition, a kind of real-time theory of composition. Again, the morbid fear of the specter "closed form" haunts the landscape.

Is an arena a closed form? Is the page? The primed surface of a canvas? (That last one sounds silly, of course.) Poets should not be bashful about borrow-

ing methods from the other arts. Why not have layered operations that affect the entirety of subsequent layers in the composition of the poem?

The idea that all writers face the same blank page at the outset is a truism. But with computers and precompositional techniques it isn't even true: one can start with a full page, and then the poet's work is to empty most of it.

EVOLUTION

The cycle: words are eaten, become compost for the next generation, become the food for pages that spring to life full, not empty. Like evolution in nature, chance may be used but is not the whole story. (Chance and mechanism both exist in nature, typically in close confines.)

Cut-up is a venerable technique. Cut-up usually means cutting up someone else. Another approach is to cut up oneself, to compose for the cut-up (precompositionally): poet as builder of the forest, the whole forest, creator of the evolution game and all of the pieces. One's words take on a different value if you know that "failed" lines will be eaten and plowed back into the next generation, that successful lines are the survivors.

OBJECTS

Many poets have written about objects. (Some who have ended up wishing they hadn't.) Computer technology changes completely the "objecthood" of words. On the computer screen, the comparison of moving words with physical manipulation of things which can be held in the hand is simply inescapable. Just as physical objects may be found in the landscape, the poet may find word objects in an arena in which they are, by whatever combination of artistic choice and algorithmic mechanism suits the poet's poetics, presented by surprise. (What an irony that chance becomes just one among many classes of algorithm, that the

pseudo-random is programmed. Random number generators have their chapter in Knuth, just like searching and sorting algorithms.)

But the meaning of the word "object" is itself changing. Object has become a technical buzzword in computing. (Such buzzwords now seem to work their way into the general language with frightening speed.) In the computer concept of object, the nounishness of the object is receding in importance; object becomes the cluster of verbs that make sense when applied to the object, with noun properties along for the ride but opaque, unobservable but through the action of verbs. The object becomes the animate object: it behaves. Whole new galaxies of animate word objects await creation.

In the animate object, juxtaposition becomes invitation: the computer may be used to define logical and physical space allowing phrases to be juxtaposed and still individually accessible; presentation of the juxtaposed elements may be part of an animate word object's behavior. This type of juxtaposition may be and or or: the poet may be inviting the reader to choose any one of the elements offered, or may be offering the cluster of all of them as a single entity, that each one and the next one and the next one are all there together. Structural zero becomes an empty container filled by participation of the reader.

Evolution: new species in the word forest, an infinity of possibilities. An arena with structure that is still open, that behaves, that invites.

BIOS/THE LOGOSPHERE/THE FINITE-MADE EVOLVER SPACE

Bios—as in biosphere, the universe, the totality of living things, the breathscape. The evolved: a generational, incremental one-to-the-next change. Contribution, frequency, probability: slow morphing of the gravity of the rules of survival. Diversity: the multiplicity of beings in the same space by happenstance. Not designed, but not undesigned: feedback. The happenstance has a reaction, self-reaction, selves reacting. It is folded in its own echoes.

Digital: the realm inside a membrane of fixed possibilities. Fixed as in the fixed number of possible bits (two) or letters of the alphabet. Or words, more or less—less because this too can evolve, does evolve. The words are there before we get there, though we can cement pieces at hand together, neologos. But we don't invent all the words, or even most of the words. At best we make a few. Mostly we choose. Selection from the fixed prior set. Finite. The dictionary fits on a CD-ROM with room to spare. The dictionary, even by the standards of the size of the operating system on a computer, is a tiny amount of data.

A subset. A selection of that infinitude of possible human experiences. From only one, or maybe a handful, of all the languages of the earth. Logosphere but specific sphere. A minds-inflated small-point universe. A set. A small set by the standards of databases, but large for the mind: one can't hold the whole language in mind at once. Robert Duncan told us—over and over—about keeping the numbers small. He insisted he could not count beyond five. Of course he was speaking about prosody, about how many choices he could keep in his mind at once. Small sets. Tiny sets.

Logosphere the model. What if we make the word set as small as possible. Not the whole dictionary but a few hundred words. Words sliced from phrases raised as domestic animals and then offered up as metabolism material, cut up, as food. Logosphere as biosphere: the energy to be word-eaten, processed, an evolutionary catalysis space. A word mass set out as the energy source for the next generation of phrase making: prompt-sheets. A glance source. And the process

cycles: a next generation of phrases written from the prompt sheets, and then cut up, permuted, eaten, metabolized: composition by evolution.

Metabolism: as in chemistry: to liberate energy you break the bonds. Cut the phrases up to get reactivity. Composition in advance just to make raw material: precomposition. To make the phrase-bonds just so they can be broken. Free association is OK as metabolism meat: when associating from A to B, keep A, keep B, don't keep the link from A to B.

On paper it is hard to see how this is achieved. You write A, you write B, they are there together on the paper. If you write them on the screen, they are there together also, but you can break the link: cut the phrase between A and B, scramble it by random permutation. A and B are flung apart by the laws of chance. The bond is not valuable, but the journey got you A and got you B. They settle into place, where they may or may not get eaten. And then, dance-wise, settle into a new place each time the permutation is rerun.

So the cycle: start with phrases, written by hand, the old way. They don't need to have meaning they only need to have energy. Cut them up. (They are made to be cut up.) Cut them in fact exactly where they mean, leaving the boundary raw, the energy pulsing out. Now take the fragments and permute them, by chance. Pour the result into a single pages long solid paragraph: the word set. The prompt sheet for the next phase in the cycle. A source. Possibly a strict source for all the words of the next phase, but not necessarily strict. You don't insist that every word for the next phase come from the set, but striving for that is an amazement: so often something at hand is right there in the available prompt sheet, a combination that works exactly in the right slot.

Extraction from a miniature totality. A multitude of topics but then cut and scrambled, thus embodying the model space, an ocean, a full range for years of work at a time. A subset language but is it so sub-: a model. An experiment. A place where the energy mingles are staged. A catalysis scaffold.

The poem starts not from an empty page, but from a full one: the prompt sheet, the source rattle, the extraction space. Its own small universe. Not read,

literally, but eye-dance-scanned at random, seeking the leap-out. Words that are there in advance of making the poem as an evolved/evolver space, a logosphere, a word space formed partly by chance and partly by building.

So where, exactly, is the digital world so different from the physical world at this business of making logospheres? We can make a logosphere physically. With no computer at all. Words as real physical objects, each on its own space. (An idea I stole quite blatantly many years ago from the painter Mary Jean Kenton.) Objecthood is not an issue: you hold the words in your fingers. Stick them to the walls, the floor, lay them loose, let them loose in the world. Let them float on water. Connect them up with drafting tape. Or as Catherine Marshall did, put them on the spheres of molecular models: word molecules. Shade them. Hide them in piles. The page becomes any possible physical space.

As physical objects the words have weight, have friction, reproducing the word means manufacture: it costs. The space even with no words: it costs. The cost occurs again for each space of words. Cost dedicated to that specificity of word. Cost not for tools or equipment but for bare raw materials of the word space.

Friction, the resistance, the difficulty of motion. In the digital world we can reproduce the words, the word space, with no friction: there is no manufacture. Propagation of the word is free, the reader can do this.

But we pay: the viewport is tiny. Poets can't afford a monitor the size of a wall. Thus the paradox: our virtual page is theoretically unbounded yet often physically smaller than the real page. So there is this motion problem: motion of the viewport through the text. A retina that is way too obtrusive.

Non-specificity: the word-set, the reservoir is not specific but was made from what once was specific. Specificities accumulated, generality by having lots of them, and then the specificity is intentionally broken up, the pieces scattered, mingled. A kind of specificity in the word-gene not the word experience: specificity is in where the word was, not where the word is. But the word space cut for the

scrambling is not cut at every word boundary. There are shards of the specificity left.

Inclusion: logosphere as a miniature universe needs to include as much as possible. But Shannon tells us: the information measure is based on how much is excluded. If everything is there, nothing means. The Shannon Measure of information is based on the probability of the codon occurring in the code: if it's always going to be there there are zero bits. Thus Shannon measures by exclusion: information measured by how much is not there. A high number of bits comes from a low probability. A paradox: one wants to maximize the number and kinds of energy transaction that can happen: to include them. But Shannon tells us that for information we have to exclude.

The specificity, the Shannon exclusion, is time shifted. It happens in the future, in the mind of the reader. It happens in the past: the word materials evolved from pieces that had some specificity but the pieces were then metabolized: getting at the logos as a precompositional evolution movement. Logos as an event horizon with a past as evolution meat and a future as the reader's energy transaction but absolutely no present. There is no present. There is only a moving blade between that specificity of the past when logos was encased in a rigid body and that specificity of the future where the energy transaction happens. The present is a phase-change boundary. Word-energies flip. Logos not as the code but as the channel. This is how we evade the Shannon Paradox. The text is not the code, it is the channel. We turn Shannon inside out.

There is no mapping between precompositional time and compositional time: The chance of a mapping is destroyed, cut up and permuted out of existence. We are familiar with earlier forms of poetics which are real time poetics—e.g. projective verse time sequences acting as the image of the poet's breath-time at the moment of composition. Or Allen Ginsberg's Improvised Poetics: the poem composed live into a tape recorder as an improvisational act, the whole act of composition being a real-time phenomenon. But why should the poet's time matter? Why not the reader's time?

Logosphere as process: the evolver space. The codes create channels which create codes which create channels: a self-sustaining energy matrix. Wittgenstein taught us about language games, but every word is a game, word is nothing but game. It flips, like the duck-rabbit. Channel/code/channel/code: there is no real identity as channel or code, it is both, it is the oscillation. It is the whole uncollapsed possibility space, as in quantum mechanics. The word as cloud of possible arrivals.

The logosphere is a biosphere: logos/bios: they are all the same. The word as a being. It behaves. Its fate is to be eaten.

QUESTIONS ABOUT THE SECOND MOVE

The first move being what it has always been—hearing the words. Or seeing the words; whatever way they first come to mind. Or perhaps they are not even words but fragments of words, or even letters. They come to mind and then they must be put somewhere: The Second Move. Of course the traditional answer to the question of where the words go in the second move is: into a notebook. Or even perhaps right into place in the poem itself. The second move is so habitual—subliminal almost—that most poets would react quizzically to even giving it a name, or thinking of it as a move at all: it is as natural as breathing. (Perhaps it is breathing (Ginsberg 1971, Olson, 1950).) But for the cybertext poet, this is suddenly an issue, not simple at all. What is the cybertextual equivalent of a notebook? Where do the words go when they are still only a scrap? Must the cybertext poet give up scrap collecting?

Philippe Bootz describes a vastly different approach. The cybertext author creates a texte auteur, a kind of scenario by which the cybertext is assembled. The cybertext author becomes like a filmmaker, having to extrapolate in the mind how the final piece will work. There can be scraps of texte auteur material—but these are not pieces of cybertext, they are pieces of directions for assembly of cybertext. The poet collects not scraps but "metascraps." Despite all of our vaunted rhetoric about the interchangeability of reading and writing (e.g. Joyce 1995), despite all of our emphasis on the interactivity of the cybertext, the writing process is in fact not interactive at all. There is no immediate feedback between the author and what replaces "the page" at the level of scraps—at the point of the Second Move—no ability to react instantaneously to those nascent bits of word behavior and say, no, that's not quite right. The texte auteur is a one-way street.

The question is not how to go from the Second Move to the Third Move, but how to make a Second Move which is (instantly, on the spot) returnable to the First Move. A Second Move which is already cybertext. A Second Move which is already interactive. How can we say we know what interactive writing is until we have done this?

So where is the code in this second move? Where is the algorithm? And can we say we have cybertext with no algorithm? Surely there is no code at all in the First Move. The poet "hears words." The only code is the code of the neurons. (We do still have the term "natural language"—as opposed to "formal language"—for a reason.) Must the code be written on the fly as part of the Second Move? How do we put the code into the cybertext poet's fingers? For the Second Move is a matter of fingers. If the First Move is a matter of ears, of eyes, of breath, it is at the Second Move where something becomes made, some object which exists somewhere for the first time outside of the poet's head. A thing has to happen. Never mind that it might be a digital thing. It is still an object, something preservable. It may be the Second Move but it is the First Object. The note. The scribble. The scrap. Something tangible. Something the poet can forget about without losing it, and have the scrap to bring back to mind. But a cybertext scrap. A word object which (already) behaves.

The question is one of arena. Where does the word object go when it is fresh, instantaneous, just born? The usual answer to this question is that it goes into some kind of cybertext authoring system. Like Flash (Macromedia 1995), for instance. But will it play at the Second Move? Perhaps not. Perhaps it needs a whole support structure, that has to be provided along with the words. Perhaps code needs to be written first. Or backgrounds need to be made: other layers. The word at the Second Move is a vagabond with no camping equipment: it has no place to go. So it fills out a form, sits on the sidewalk, waits for the builders to arrive. This is what the texte auteur is: a requisition. A request for service. A request for builders, with a list of things they need to do. In my own case, not even a word object at all, but a notation in a kind of intermediate language, e.g. "Build a word cluster here with the following layers ...". The word has to wait, perhaps weeks, before the word object gets built, according to requisition.

So then, what is to be done, to allow a Second Move that will play, that doesn't need the requisition form? The poet needs something to act as a notebook. As the first wet-paint home for scraps. Not just word scraps, but word scraps that behave. That can move or respond or be layered—whatever kind of behavior the poet's aesthetic calls for. Behaviors that are pre-built, receptacles waiting for the

words. No code needed at the Second Move because the code is already there. A construction kit for word objects.

Of course environments like Flash already provide an assortment of prebuilt objects. So what exactly is the problem? There are several. At the moment of the Second Move, the final home of the word object scrap is: Don't Know. Just as for centuries the final home of a line jotted down in a notebook is: don't know. If writers have been doing this for centuries, why does software make it so difficult? The corollary: at some perhaps much later move, the poet does know: the word object does get a home. So it should be very simple to just move it, to plop it into place. With all of its interactivity or animation or behavior of whatever kind coming right along. Just as the syntax comes along with the words when you paste in a sentence somewhere. (Syntax is a kind of behavior. It tells how the supposedly linear word stream folds and unfolds, how parts of the word stream attach at a distance. How pieces of the text attract or repel. Places where the sockets are, other places that go into the sockets. Almost a kind of animation as the connection happens.) Cut and paste is so simple to use in an ordinary word processor—at some later move the results of the Second Move must be pasteable—somewhere else.

The set of kinds of objects available in Flash at the time of the Second Move is closed: you can't extend it. You can create new objects for the reader: objects that work at run-time. But what about new objects for the writer? New kinds of objects that can be used with the poet's fingers. Just after hearing the words. With no sound of code to drive away the words. Just like drawing: a swoop, a hand gesture, a motion. Clay being layered up, oozing out from the fingers as the wheel spins. The tactile feel of the words being molded into place.

And then: you play it. Immediately, on the spot. The Third Move is to play the word object fresh and instantaneous, cock your head, wonder: how do I like how this plays. Is it right. Not a question you have to wait weeks to ask, but a question you can ask on the spot. Just after having heard the words. A Third Move that happens so soon after the Second Move, perhaps the Second Move took so little time you aren't even aware of having made it. A second move that is disappearable, out of mind, like breathing. From the First Move to the Third Move as

if there never was a Second Move, as if the Second Move is built in to the First Move. Second Move: go away! Recede, become subliminal.

The requisition still happens—but this requisition is for the programmer to create new kinds of objects in advance. In advance of even the First Move. This is not a texte auteur for a specific piece, but a requisition for enlarging the construction kit.

So how do we do this? How do we specify the kinds of word objects we will need? We need to become almost more like sculptors than poets: we need to send an order to the foundry for fabrication of new kinds of word objects. But here is exactly the difficulty: there is no foundry. We have to do it ourselves.

There must be a system of objects. Of movable objects—movable with your fingers. Objects that allow the poet to make new kinds of objects and plug them in. Objects like an attractor that would pull in words or phrases or letters from a surrounding text space. Or a repeller that would send them away. Kinds of objects. (Programmers call these "classes.") A useful variety of kinds need to be available in advance, but also the poet should be able to add more, with reasonable effort. You will say my arithmetic is wrong: if the poet adds a new kind of object to the system, what move number is this? It happens prior to the First Move. Is it a negative-numbered move? It is a move from a precomposition layer, a layer that may only need to be done once to support a whole series of pieces: like building a desk. Not the writing but the carpentry. Trans-writing. A producing, like the producing the programmers did who wrote the Flash development environment except that this one is open, allows us to add or own kinds of production alongside the ones we were handed. Open to new objects not when the piece is played but when it is written: at the point of stage and scene and frame. The poet must be able to change the way the stage works—to be one's own stagehand.

This is not a pipe-dream. There is such a system. Squeak (Ingalls 1997, Guzdial 2001, see also www.squeak.org) for instance. A system that can support words as graphical objects, animations, new kinds of objects invented by the poet, objects that "carry" other objects around with them. A system that lets you move an object where you want it, drop it into place, and all the behavior (by itself and

by its components) comes right along with it. A system that is available for free (including source code). A place to start.

An arena. An object arena. A place where small active objects (e.g. animations) can be put into a sentence as if they were words. A form of interactive writing that lets the Third Move be playing the results of the Second Move—immediately, on the spot. The way the poet has always been able to read the line just after the Second Move—while the ink is still wet. Wetness: our cybertext composition process needs to become wetter, like wet clay. More liquid. Or like a gel, spreadable with the fingers. The place the breath breathes into. The potter's wheel. The clay tablet, a still flexible digital surface.

AN INTERVIEW ON POETICS

Sandy Baldwin Interviewing Jim Rosenberg, February 2003

Sandy Baldwin: We were talking about your recent essay, "Questions About The Second Move" <http://www.well.com/user/jer/q2m.html>, which you describe as "about the necessity and difficulties of having an authoring environment where the poet can collect scraps: a notebook." You were saying…

Jim Rosenberg: It's almost like an artist studio snapshot of what I am engaged in right now, and it's really very exciting because I have given a lot of talks over the years about tools and the inadequacy of the tools that I had to work with, and at hypertext meetings the systems people were very eager to find out from us writers what our requirements were for tools. So they pushed us pretty hard to create requirements documents and I have done that. I had one published and stated explicitly what my requirements were. It was always this sort of pipedream to think that I could ever see it happen—to actually be able to make it happen myself is so exciting. It's really very, very wonderful. But you know when I see so many works being done in tools like Flash. . . I'm not a very big fan of Flash. I'm certainly not trying to disparage any of the works done in Flash, but, if you think about it, Flash was really a tool created for advertising, and the niche for Flash was that it was actually created for animations to go in web ads. Now, why should poets have to deal with a tool that was made for advertisers! What would a poet's tool look like? And what is the new media equivalent of the notebook? That's the question I keep coming back to: how do I make scraps?

S: You need a printed notebook.

J: Right, a bound notebook. I don't think I know a single poet, in one way or another, that doesn't use a notebook or accumulate scraps in some fashion or another. That's what is really hard to do in new media. I have been very concerned for a long time about the pieces that people have been doing— have done—in a medium where you really can't gather scraps properly, and can't really mold it fluidly in the same way that people are used to on paper. And what would a poet's authoring system look like?

S: Would you say that what we end up seeing with Flash poetry is the scraps you get from elsewhere, things pre-made in the system already in Flash?

J: Well, Flash doesn't lend itself to collecting scraps. I mean it lends itself to making pieces very laboriously, but the question is how do you collect fragments where you don't know where it is going to go.

S: You say that the status of the fragment in a notebook is something like "don't know." And that's what you're alluding to here.

J: Exactly. You don't know where it's going to go; you don't know where it's going to end up. I mean, for centuries poets have collected lines in various little ways without knowing what kind of poem they are going to end up in, and maybe they would go years before that fragment got used, but it's in a notebook somewhere. It's important to capture it because if you don't capture it, it might be gone. Well, how do we do that with new media content? Now, Flash, I suppose, has libraries, and people would argue you could sort of do it this way, but it has got to be almost to the point where it's subliminal if it is going to work for poets.

S: It's a granularity problem too, right? 'Cause the Flash—

J: Exactly.

S: Isn't this a question about the fundamental nature of hypertext in writing, too? Wasn't this what Vannevar Bush wanted to do with the Memex? <http://www.theatlantic.com/unbound/flashbks/computer/bushf.htm> Do you remember this? He wanted the same thing, he wanted a machine that would gather scraps.

J: I guess you could link that in. He was concerned with a sort of more limited range of behaviors on the part of the word than we are these days—

S: He certainly wasn't thinking of poetry.

J: I tend to think of Bush more as the pioneer of explaining what the corpus would look like. Particularly the corpus of associating lots of people's work together in a single body as opposed to the process of composition.

S: OK. Well, I guess the reason I was thinking that was because, when we talked about gathering scraps, I also thought about how we gather information as we use electronic media as we surf the web, for example, and how there's really not adequate ways to footnote what we find, to save what we find…

J: Right, well, I've written about this. It's called "gathering," and the state of our gathering interfaces is just awful.

S: I have a question. Do you think of Amazon as a gathering interface? Maybe this sounds facetious, but it strikes me that for many people, that's the model for information gathering, with its information cookies, and its "eight other people bought this, you might want to take a buy it." Have you looked at this?

J: Well, I buy a lot from Amazon, and no, I don't think of it that way.

S: Well, do you think that may be a kind of cheapened model of what we want?

J: The cheap model of gathering is bookmarks, and it's just not a good model. Now, I must confess that I personally don't use bookmarks. I mean, I keep an HTML page in which I create links for something I want to bookmark. I have it as an HTML page which I can load, and that way I can organize it using standard HTML authoring tools. But I used to use something called WebSquirrel, which was wonderful, and then I got off the Macintosh. I guess I could use Web Squirrel again, but I got used to using just an HTML page which is easy if you're in various types of other platforms. But bookmarking is very crude. How do you bookmark a sentence from a web page? Supposedly, the XML revolution is supposed to fix this, but I actually haven't seen much that advances the state of the art gathering interfaces outside of the research community.

S: You mean research community as in Aquanet, and things like that. I was being facetious when I was talking about Amazon, but don't you think this has a

lot to do with our freedom and when we use these technologies? If we don't come up with gathering interfaces, then someone else is going to dictate how we use—

J: Oh, I see what you're getting at. Yeah, that's a very good point, that's an excellent point.

S: Amazon sort of purports to think for us. Or I'm thinking of Microsoft's Paladium, right?

J: Let's not get started on that, you'll run out of tape.

S: But, I mean, the stakes are high…

J: Yes, and, though you mentioned Paladium, a more direct cross reference there is SmartTags.

S: Right.

J: I think Smart Tags will come back, unfortunately. Now, Smart Tags was very insidious. That was the effort to pre-gather and to just to say, well, Microsoft will do all of the gathering for you, thank you.

S: Yeah, I guess that's what I was thinking. And I thought about this in terms of where you talk about, where you talk about Flash, when I took you to be saying that you need to be able to have an artistic construction kit that will let us be able to name and code the behaviors we have at that level.

J: That's right. One of the things that I find about Flash that I find the least acceptable is that there is this very strict segregation between two vastly different models. I mean, when you really talk about the technology of Flash, you're not really talking about a technology, you're talking about two technologies, and there is this almost schizoid division between the two in which they're totally separate beasts. So there is on the one hand Flash the authoring environment—

S: Right.

J: —and then there's the runtime environment.

S: Yeah.

J: And, they're totally different. I find this unacceptable. Just as I've been interested for years and years and years in reducing the granularity of hypertext so that it's inside the sentence—morphemic is the word Cathy Marshall used for my work, which I like a whole lot. It seems to me—and I've had some direct experience with this, just in the last few weeks where I finally got my Squeak stuff working as an authoring environment—the question of authoring behavior versus runtime behavior should be completely granular at the level of the object. It's hard to talk about objects in the technical computer science sense, and really explain this to people who have not been through it. It's a radically different paradigm. Languages like Smalltalk give you a totally different perspective on the world than other kinds of programming systems. One of the aspects of this is that it's a sort of seamless environment in which there aren't these kinds of boundaries. So for instance, people have accepted to an appalling degree the point of view of companies like Microsoft on what is the division between system level software and the kind of software you can write for yourself. In a Smalltalk environment there is no division. You get something called an image and it comes with full source code and you can change anything about the system, and you don't just write a program, you subclass something and your code gets integrated into this system in a way that really—there's no boundary between your code and the rest of the system.

S: I see.

J: And everything is an object and… everything is some behavior that can be modified, it's all exposed. It's a radically different concept than this idea that… um… if you make a new media piece in Flash, you have objects that you've assembled yourself. You have certain types of scripting that only you can do at runtime, but meanwhile you are doing that within an authoring system that was written by Macromedia where you can't modify anything and that runs on top of an operating system that you get from Microsoft and whatever it does you can't

change at all except in these very constrained ways that they give you. Um, this is a very different framework from a Smalltalk environment in which the entire source code is available and it runs fine on open source operating systems where you can run it on a box where you've got a hundred percent of all the code.

S: So… So what you're writing, you know I'm not going to be able to say this right, it integrates itself right into the system it's a part of right?

J: That's right, that's right.

S: And, and you know, this was one of the things I found fascinating in the.. the essay "Questions on the Second Move," because it seemed to me ultimately the second move and first move part kind of fell together, right? Because…. And there ends up…

J: Right, it's almost like… It's almost like the idea is to abolish the second move.

S: Right, and as you say… it's a question of granularity to some degree and that's why I thought of your earlier poetics essay "Openings: The Connection Direct," because that starts with that wonderful paragraph on what you called "energy transactions," which is hard to pin down but it has something to do with… I thought of this word in information theory "phatic communication," p-h-a-t-i-c but it's not quite that, but its something that is not purely related to the idea of transmission or possession, and I may be condensing things together here, but I thought one of the interesting things about that essay, you were saying, you know, poets are concerned with this, the question of the energy transaction layer and in a way there's the same question in the "Second Move" essay of abolishing the second move, and saying we need to hold on to the first move, and I agree these are different essays and they are from fifteen years apart or something, but for me I guess I saw certain continuities in them.

J: That's interesting. I hadn't thought of it that way. I still believe very passionately in this kind of energy transaction and I still believe very strongly in the viability of composing directly to the energy transaction layer, as opposed

to the more stereotypical attitude people have about communication transmitting a thing, or an entity, or an idea, and it's like getting moved from one brain to another as a sort of infusion and both pieces are really strongly about objects and this objecthood is really something very very essential. I remember being at some event—I'm trying to remember what it was, I think it was a question and answer period that John Cage gave after, I believe it was the New York premier of Harpsichord. It's been a number of years now and it might have been a different event, but there was a question and answer period and Lucas Foss I believe organized it. He started it out by repeating to John a comment that he had heard from some disgruntled woman who was there, who said about the piece, "It doesn't care" and he handed that to John Cage, just like that for his comment, and in his completely inimitable fashion without batting an eye or missing a beat Cage got his usual Cage childlike smile on his face, and he said, "I would think the same thing could be said of a sunset." Now, my paltry scratchings have put it much less powerfully than John did on that particular occasion, and I don't want to try to become an interpreter for John Cage, obviously he speaks well enough for himself, but what I took that comment to mean was that in effect, the sunset is an object of the world, it has no author, unless you want to start preaching in religious terms and personifying a deity that way. It has no human author, it has no author in the usual artistic or aesthetic sense. It's an object, and the viewer relates to that object and there is an energy transaction there even though there's no expression, no author, there's no transmission, and it's between the viewer and the sunset. In this case it's between a person and the sunset, or it could just be a collection of stones you find that are incredibly appealing or a seashell that you find on the shore that's striking, and it's that particular object and it's your relationship the "receiver" to that object. Now, we need that kind of intimate relationship in the authoring system. It needs to be a very seamless relationship between the authoring system and playing it; it should play immediately. I mean, what does it mean that we have got ourselves roped into the kind of composition where it's only after weeks of work that you can read it. I mean, you know, it's been part of the act of composing poetry for millennia that the minute you compose a line, it's there for you to read, or hear, or play back in your head and say, "Gee, does this work here or not?"

S: And if I may, about what you were saying about objects, and this strikes me as why you... I guess here I'm thinking about an interview that was done once

with you by Judy Malloy, where there's an exchange where she says how she had always thought that in your work the user brings the poem into being by mousing over it, and you said that wasn't quite right, because you felt in a fundamental way, at some level, the objects in a poem already exist. I may be jumping too far in thinking of hypertext as something that is produced by the reader and am I going down the wrong track to say that the question of the object is something that one could hold against the reader for producing the text, that the text doesn't have a real existence.

J: Well, I'd have to go back and remember what I said in that earlier interview, but I don't want to speak against the concept that the reader constructs something, and that's really what's behind the whole concern for gathering. It's that the reader needs explicit tools to construct something.

S: Can I say this—I guess I'm hung up on this idea of the energy transaction level—as a reader, what I construct... if I think I'm just getting some information from you, that's not really the point, right, I'm supposed to perceive these objects at some level.

J: Well, we're getting into another place, which I feel very passionately about, which is very dangerous, I need to be very emphatic about this. I don't want to say anything that implies in any way whatsoever that the reader has any obligation at all! I don't know what it's like for kids in school these days. We don't have kids—I don't know what textbooks are like at the elementary and high school levels, but I know what they were like when I was a kid and I know how poetry was presented, and poetry was always presented being surrounded by this sticky gooey mess of obligations. There were all these things that you were supposed to do—

S: And supposed to get out

J: —of a poem. If the result of postmodernist theory of the reader constructing the meaning is to add another layer to this level of obligation on the part of the reader, I'm gonna say, "No!" The reader is under no obligation whatsoever. What I want kids to know about poetry is that a poem is something you pick up

for the sheer hell of it and nothing else. So, one of the things I think is important to understand about the energy transaction is that is isn't necessarily something conscious. In fact, it could be something anti-conscious.

S: mmhmm…

J: You could have a conscious impression while looking at a work that is completely awful, that you are getting nothing from it, that it is just horrible in every way, that you just don't relate to it, and yet there are ways in which underneath the surface, subconsciously, you could still be receiving some sort of energy from it that could have a positive impact.

S: So in this way we can't really, well it's beside the point to talk about whether it was a good energy transaction or not. Now that's not what we're talking about.

J: Oh well, I don't know. I don't know about that either, I mean think about food. We all know that we have food that is delicious, the type of food you brought, and food that we can't stand… and we also all know that the nutritional value of food is not necessarily related to our pleasure in eating it, and you can have a meal you thought was delicious and then you know, you just… particularly if it's a meal you ate in the beginning or the middle of the day, and it takes you all day to realize that it wasn't good food as fuel. And that is really the issue to me as far as the energy transaction… Was it fuel? And I wouldn't go as far as to say we can't talk about that. That seems going a little far… Of course you might never be able to put into words what it is that makes something good fuel or not.

S: I know this is not a direction you look in your writings, but I… It makes me think of the phenomenological philosophy, you know whether it's a question of Husserl or Merleau-Ponty.

J: Yeah, well, I've never really read those people, unfortunately.

S: Can I ask then… I don't want to leave this, because this is quite interesting… I was curious about the role of the visual image or the visual in your

poetry, because there were a couple of questions I had in that direction. Maybe the first one is just a simple one: What do you think of the relationship between your poetry and concrete poetry? And, maybe more generally, people think of themselves as more visual poets because they are concerned maybe with some of these issues…

J: Well this gets tricky, I really want to be careful because I don't want to speak against the type of work I don't see myself as doing. Umm… I should say…. and I feel strongly about this… If you would ask me "What adjective do you want in front of the word poet, in reference to yourself?" Um, I would not disown terms like "new media poet," or "digital poet," or "hypertext poet," or what have you… but the adjective that I would really prefer is none at all. I see the work that I do as simply poetry. Now I do use visual means, and the way that the work looks matters. I do spend time on it… I certainly don't think of myself as a visual artist, I don't do other kinds of visual work. I started using visual means because it was the only way I knew how to achieve the result of putting words on top of one another and having something readable, and putting words on top of one another and having something that could then be put into a larger unit the same way that a word can. So there are these visual means being employed that are not just words. There's a visual notation and so forth. But I'm not doing this to supplant the level of things like syntax, semantics and so forth. Some visual poets have been very ideological about this and so it's got to be instantly accessible to someone who speaks any type of language, and so forth… Um, It hasn't been the way I've wanted to work at all… so the visuality is simply there as an added channel… The media allows for this possibility so I'm using it.

S: But isn't it, isn't it…a grammatic quality of the visual as old as your poetry to the degree that the new media poems automate the diagram, right?

J: Right…

S: You've been doing almost the same visual task but also working on some of the same visual problems for 34 years. Not just the poetry has a visual aspect but also a visualizing aspect because the visual aspect the diagram is something

that viewer has to interact with, the reader. It's not just a picture to be looked at is what I'm saying…

J: Right, right…

S: I don't know if I understand exactly what the reader is doing, but this goes against the idea that a picture is something over there that I gaze upon…

J: Right, right.

S: Do you think, is there, are there sentences that can't be diagrammed?

J: Well that's a provocative question, that's the type of question I would have to think about and then two years later I would probably have the right answer. But, umm… The the concept of sentence as we know it is inherently bound up in the idea that the words have structure, that they are combined in a structural way. And the diagram is only one of way of illustrating that structure. And what intrigues me about the diagram is that it is a way of opening up the structure to words that are combined with no structure. I've said this before in some of my other writings, that Cage has once criticized the twelve tone system as not having a zero. And I am not up enough on twelve tone music theory to know whether he was right about that or not, but um… syntax typically assumes that every element has a structural role, and the thing that syntax won't do is allow elements together that don't have a structural relationship. So I guess I would turn your question backward, and say that the interesting thing to me is not whether every sentence can be diagramed but whether it is possible to include something which cannot be diagramed into a sentence which can.

S: Then what would we have? Do you feel like you approach that? Does that happen?

J: Sure, sure… The word clusters are just, words are sort of just there together, like two sounds that are played together, or two visual images that are superimposed on a canvas. Um, but yet with the diagrams, I can include those things into a sentence as if they were words.

S: It strikes me that that comes close to some kind of traditional connotation of poetry as the new. This is something new, some new innovation, right?

J: Well, it's hard to talk about this without making it sound like you are creating some advertisement, and I will readily confess that I fail at self promotion, that is one of my great failures. Umm, I guess this comes back to our earlier discussion of energy transaction. I have never felt comfortable with this sort of blind allegiance among certain people to a banner of "making it new." I'm uncomfortable with that. It sounds too much to me like our stereotype of manufacturers who make, build, products with planned obsolescence.

S: Yeah…

J: I mean, It's just problematic every which way, I mean the, the issue to me is not whether it's new, it's whether it really has crystallized that energy which has been available. Is it energetic? That to me is the more important question than is it new?

S: Umm, It strikes me that the really interesting point that you made a few minutes ago about, the possibility, if I get this right, of diagramming or bringing into a diagram something that doesn't have a structural relationship has a certain similarity to some of the questions you've considered in types of hypertext, for instance the possibility of a null choice.

J: Right…

S: You know, and I have two different questions here, and we can go with either one. One: Whether you'd want to talk about the null choice in hypertext… or the other, more general question about the relation between what we've been talking about here…which is your poetics (on the one hand) and then this other mode you write in, the more academic/scientific mode, which is maybe separate and maybe the same, where you consider genres of hypertext and hypertext systems. So one question could be what could you say about the null choice and

hypertext and the other, are there different flavors of Jim Rosenberg's writing here or not?

J: Oh dear, I don't know where to go with this, umm… I don't remember if I ever exactly used the word "null choice" but maybe I did… umm, "Navigating Nowhere" is an essay that's about…

S: Yeah, that's what I was thinking of…

J: Yeah, okay that's what you were referring to... Maybe I did use the term null choice in there, I know I used the term "no structure." Umm, this gets into some very tricky stuff that is absolutely related to discussions of the craft of poetry that have gone on for centuries. Lots and lots and lots and lots of people have written lots and lots and lots and lots of words about what constitutes a good line, and what is the role of the poetic line and if you extract the line from a poem and just look at it by itself, sort of, what is this creature that's pulled out from its environment. And we really don't have, I don't think, a good body of theory for this in new media poetry. We've sort of evaded these types of questions… so if we were using the traditional, what I like to say, disjunctive, type of hypertext theory, node-link hypertext… the question would be, "What is a good lexia?" What actually is the role of the lexia? What is the role of the screen that you just get when you uh, go nowhere… Now there is some kind of new media poetry where that question doesn't occur because you have animations that play and you don't have the choice to do nothing, and you can't stop it. You could stop it by hitting escape and quitting out of the piece completely. But you can't freeze it.

S: So we are near to the question of…

J: But the reason that these things are tricky is that one does not want to be prescriptive I don't think. I feel really strongly when people ask me, "what is your definition of hypertext," I feel obligated to say, my favorite answer is, "we're working on that." I mean this stuff is so new that I think it is a little soon to be prescribing "Oh you have to write a lexia this way…" and this is "how you do a web page for usability" and "you should do this," or "you must do that"… it just gives me a headache. But nevertheless, I wonder sometimes whether we've evaded our re-

sponsibility as hypertext writers at really dealing with some of these same types of questions of craft and poetics that earlier generations have dealt with. The phrase that I like to use is that the various pieces of a hypertext have to hold their aesthetic weight, pull their aesthetic weight… they have to hold the load. Now what that is could be subject to differences of opinion, and people will have different aesthetic points of view on this… so the lexia even if you assume that, let's take a hypertext piece where the emphasis is on the links, or maybe the emphasis is on the links to the degree that they have these timeouts that will force the link to be taken if they've been sitting there for a certain period of time. And the author has put all this effort into the links but the lexia is there, and it plays some role. Even if the author is saying to me: "look, I'm not interested in what the lexia does, I'm interested in what the links do," we can still say okay, that is a valid point of view, but if you have an unhappy lexia, it could get in the way of that… It can get in the way of what you want to do… so that question is there. And this gets back into the whole issue of tools. Somebody composed that lexia… Were links being followed when that lexia was composed? Now at some level there is something there in the hands of somebody where it's not moving, or it's contained in some sense. And it is a legitimate question to ask at some point, "What is that aesthetic?" "How does it work?" "What is the relationship between the parts that are fixed and the parts that move or the parts that ask the reader to move to some other place?"

S: So is part of it that there is a confusion of substance…where people think, for example you're looking at "Is my lexia alright If my links are bad" and vice-versa. Whereas you really don't have one without the other. I mean I take a certain amount of what some of your essays do to have almost a rhetorical function when you ask these questions, I think this is a question that you've repeated… "What's in a link?" "What is it?" Sometimes I get a sense that you are trying to force the reader to, I guess I'm going to say, change granularity, change the level that they're looking at these terms.

J: Umm… Well it just all goes back to this idea that you've made something and you've made it and you are aesthetically responsible for the result. There are some painters who make canvases with some very unusual shapes. In this type of painting the stretcher itself is something that the painter wants you to pay attention to. Now there are lots of other types of paintings where the stretcher is just

there as part of the framework and it doesn't show… It's inherently behind… it's part of the infrastructure but you don't really see it and the painter may not want you to pay attention to the stretcher. But if you have a badly done stretcher or you have a stretcher that's fine and then warps, then you've got this funny thing sticking out of the corner and you do see it and that's a misadventure, so we can't walk away from these aspects. If you make a new media piece, if you chose this and you chose that, and it's all part of a piece…I'm not sure where we're going with this…

S: Well, do you think the answer is self-consciousness on the part of the new media author, or self awareness of choices made? Or do you think we need… If on the one hand, we're not served by the flash, the flash of the pen, and we're not served by prescription either, is an answer, kind of, pragmatism, where people perceive thoughtfully?

J: I guess all I'm trying to say is, and I don't want to labor this point too much, …when I get together with other new media poets and writers, it seems like we spend almost all of our time talking about "techie" things. You know the aesthetic questions in this type of writing as in any other type of writing. The relationship between the technical means and the elements themselves, the words of the piece, be they words or visual elements—whatever they might be. This is where the whole "making it new" thing, and the word "new" in the "new media" is in some ways really an obstacle. We seem to be so caught up on this, I'm not sure how to say this in English, techne…we're so focused on that…I sometimes get the feeling that we're too eager to let those issues be something we can hide behind, but it's just all poetry really. I mean the dot com phenomenon turned out not, after all, to be a new economy. The laws of economic gravity were, lo and behold, not repealed, and just as there wasn't really any new economy, ever, at all—just this great speculation, you could make the case that there really are no new media. What we have are some interesting things in the tool kit that we didn't have before, but, it's poetry.

S: Right. Thank you!

A CONVERSATION WITH JIM ROSENBERG

on the Interactive Art Conference on Arts Wire

January 1996

Judy Malloy

Poet Jim Rosenberg is our January guest on the Interactive Art Conference. In his own words:

Judy Malloy

Jim's Intergrams - layered poems read in part by "peeling" off layers of text - and his idea of taking hypertext into the very structure of language are radical/individual/visionary ways of looking at working with words and computers.

Jim, can you tell us a little about how you evolved (if that is the correct word here) from mathematician/computer scientist to poet? Has this background been influential in the development of your poetry?

Anna Couey

Welcome, Jim! I'm looking forward to this...

Jim Rosenberg

It wasn't so much an evolution as a "co-occurrence" of both sides to my life all along. I've been writing poetry since 1962, and have been programming computers since 1964. I guess I must have started playing with mathematics on my own outside of school back around 1960 or 1961. The two sides haven't always gotten along well. In the summer of '68 I decided computers were screwing up my head for writing, and resolved that after graduating and after my summer job was done the next summer, I would throw the computer out of my life completely. I did actually do this, and went to grad school at Berkeley in math; after the summer of '69 I don't think I touched a computer for 8 years. My trajectory as a graduate student in math ground to a halt in 1974 [long story] and for a few years I tried living 100% as a poet.

Not only did this not work economically, I found there was something missing from my life. When the microcomputer revolution hit it allowed me to work out the problem of earning a livelihood while doing the artistic work that mattered, but still the two sides fight with each other occasionally. For the last few

years my life has seemed pretty much in balance; though holding down a regular job leaves me less than a third of the time I'd like to *read*.

There are some pretty obvious mathematical or logical aspects to my work; I've always considered that that was "just me" -- the two sides are me, after all. (This hasn't always been easy for other people to swallow; John Cage once told me, criticizing some of my earlier diagram work, that he found it "too mathematical" -- as if this were an insult. This actually hurt quite a bit, but after a time we got past that point.)

I guess I would have to summarize by saying that my appetite for things computational and mathematical on the one hand, and for poetry on the other, have led me to pursue both, since roughly the time I was in high school.

Incidentally to call me a "computer scientist" is being a bit too generous: I have absolutely no credentials in computer science (though I do have *a lot* of work experience as a working programmer.)

Judy Malloy

Thanks Jim! When/how did these mathematical and/or individual aspects - what I refer to as your "structuring" become integral to your work?

Anna Couey

Not to deter you from answering Judy's question, because I'm very interested in what she's asking and your responses...but I'd also like to explore another track. hope you don't mind a polyphonic interview/conversation :-)

In your description of Diagrams 4, you point to the work being made "without manufacture" and published "on demand". Yet you do not term Diagrams 4 as interactive. From what I can gather from the sample of Intergrams and your subsequent work - interactivity there consists of the reader peeling off the layers?

I'm fascinated by the spatial, connection-oriented way of reading your work impels. The diagrammatic structure and the sensuality of your words (sensuality as in evoking the visual, tactile, aural) make for a reading that is both a physical or sensed experience, and one that is highly conscious - requiring activity on the part of the reader to make connections.

(& btw, it's a real treat to see your web site - & how your work has evolved over time.)

What determines, in your work, the designation of "interactive"? What role does interactivity play in the structure of your poetry.

Douglas Cohen

Greetings Jim, it's great to see you here!

This is already quite fascinating, and I don't mean to add yet a third tract to this conversation, but I will anyway.

As a composer, I'm interested in how music influences your work (the mention has already been made of the analogy between your "diagram" poems and tone clusters). (One could also pose the question as, "how your work influences music.")

Am I correct in recalling that you wrote a program for John Cage to automate his "writing through" texts (I think it was called MESOLIST)?

(Well, I guess that opens up two more topics.)

Anna Couey

More good questions :-) Glad to see you here, Doug.

Jim Rosenberg

Judy:

> When/how did these mathematical and/or individual aspects -
> what I refer to as your "structuring" become integral to your
> work?

That's been a long process. I began experimenting with my diagram notation while still an undergraduate -- back around '68, as I remember. From there through the early '70s I worked with it off and on -- I varied a lot back then in how "experimental" my work was. I guess I became convinced somewhere about 1974 that this was at the core of what I had to do. I don't have any clear memory of just how I decided this. I simply had the sensation that was where the energy was coming from, that this was the work for me to do and no one else.

Jim Rosenberg

Anna:

>> In your description of Diagrams 4, you point to the work being made "without manufacture" and published "on demand". Yet you do not term Diagrams 4 as interactive.

What determines, in your work, the designation of "interactive"? What role does interactivity play in the structure of your poetry? <<

The Diagram Poems were originally written to be seen on paper. The only interactivity -- in either the version you hosted on the ACEN conference of the WELL, or my Web site -- is choosing which one to read from a menu. That's not really much in the way of interactivity.

Somewhere along about '86 or so, playing with bit-mapped graphics and a mouse, I realized that software provided me a way of doing something I had wanted to do very much from the very start: word clusters -- putting words literally on top of one another. When words are put on top of one another visually, or aurally, the result often is that they interfere with one another to the point of unintelligibility. With interactive software, the words can be put atop one another and then by using the mouse, the reader can reveal individual layers one at a time, so all the words are intelligible. I don't think of this as peeling layers off so much as diving into the simultaneity to look at the layers one at a time, though I suppose either concept would work just fine. Interactivity also provides a way to navigate the diagram syntax when it's too complex to fit in one screenful.

With _Intergrams_, you simply can't read it at all without "operating" the poem. Incidentally, I should say that the actual user interface doesn't really work the way the interface works in the extract I have on my Web site. In the actual piece, layers in what I call a simultaneity are revealed simply by moving the mouse through invisible hot-spots -- there is no mouse-clicking, though I do use conventional "click" buttons for navigating the syntax. (There simply *is* no way to do no-click hot-spots in HTML -- though I'm eager to see whether perhaps Java will let me do this.) In the Diagram Poems you can simply read a poem printed out and there's nothing to "operate"; in my interactive work, you don't get to the words without operating interactive structures that are deeply embedded in the text.

This is taken to an even further extreme in the piece I'm doing currently. The simultaneities are nested many levels deep, and there are simultaneities in some cases inside the sentence itself.

>> I'm fascinated by the spatial, connection-oriented way of reading your work impels. <<

My work has become even more and more spatial as time has gone along. There are a lot of parallels between my work and the hypertext research carried

on by Cathy Marshall and her colleagues -- she's done probably the most important research in spatial hypertext. Reading her Aquanet paper I just about jumped out of my seat -- she was using a diagram notation for what are known as Toulmin structures to diagram argumentation that looked so much like my own diagram notation that it was uncanny. It seems as though there is being more attention paid to spatial hypertext lately than there used to be, which is encouraging.

Jim Rosenberg

Doug:

>> As a composer, I'm interested in how music influences your work <<

Music influences me constantly, and I guess it's always been that way. I get tremendous amount of energy from music, and would have to say the single greatest artistic influence on my life was John Cage. Beyond the specific idea of tone clusters -> word clusters, it's hard to be specific about what exactly turns into what in terms of music influence, but it's there, as energy absorbed and recycled. These days I find myself being moved the most by composers like Ingram Marshall and Pauline Oliveros, though I like a lot of different kinds of music.

>> (One could also pose the question as, "how your work influences music.") <<

I would be thrilled to think of having influenced music somehow, but I haven't had any composers tell me this.

>> Am I correct in recalling that you wrote a program for John Cage to automate his "writing through" texts (I think it was called MESOLIST)? <<

Yes, I wrote some of the early programs in Cage's mesostic projects. Andrew Culver ended up doing most of Cage's programming.

Judy Malloy

> along. There are a lot of parallels between my work and the

> hypertext research carried on by Cathy Marshall and her

> colleagues -- she's done probably the most important research in

Thanks, Jim, for all your responses. I've been working with Cathy Marshall on a collaborative hypertext and in addition to the uncanny way that our voices blend, it has been good to work with someone whose views about what is and what is not hypertext are so open.

> so much like my own diagram notation that it was uncanny. It
> seems as though there is being more attention paid to spatial
> hypertext lately than there used to be, which is encouraging.

You and I have talked for about how difficult it is to work with experimental structures in a hypertext community where work is criticized because it is not hypertext. I've had acceptance problems with my work (like Wasting Time that Richard Gess published but no one else has been interested in) that is based on parallel streams of text. Originally I didn't call my central work (Uncle Roger, Penelope, The Yellow Bowl) hypertext but narrabase (derived from narrative database) I still like that term but found that Mark Bernstein was correct that it is better to be associated with the dominant trend. Actually talking to Cathy helped me see my work more in hypertext terms.

I was blown away when I first put mouse to your intergrams but I didn't think of the work in hypertext terms, but if you look at hypertext in terms of complete building blocks of information put together in various ways - which is how I now look at it (but perhaps you see this differently?) than it is hypertext - spatial hypertext. Are you comfortable looking at your work in this way/entirely in this way? Does it matter to you what it is called?

And to throw in another loaded question, can you talk about language poetry and your relationship to the language poets?

Jim Rosenberg

>> Thanks, Jim, for all your responses. <<

Thanks all for the questions!

Judy:

>> I've been working with Cathy Marshall on a collaborative hypertext and in addition to the uncanny way that our voices blend, it has been good to work with someone whose views about what is and what is not hypertext are so open.

You and I have talked for about how difficult it is to work with experimental structures in a hypertext community where work is criticized because it is not hypertext. <<

Well, maybe I'm being naive, but it seems to me there is reason for optimism. The conventional "node-link" model of hypertext is still dominant -- and the Web hasn't helped that any, though Java may turn things completely on their head in that regard. But it seems to me the hypertext research community *is* opening up to a wider concept of what constitutes hypertext. I'm encouraged, for instance, to see that there will be a whole session at the upcoming Hypertext '96 conference in March devoted to Spatial Hypertext. (With a certain very curly-haired character named yours truly one of the speakers ... :-)) Eastgate has a new product called Web Squirrel openly advertised as spatial hypertext -- the first really commercial spatial hypertext product I'm aware of, actually. (Mark Bernstein openly credits Cathy Marshall's VIKI as the origin of many of the ideas for Web Squirrel, BTW.) I can't say I've experienced the kind of "discrimination" -- if that's the right word -- you have, Judy. That may be because I'm working in poetry, and there is so much less hypertext poetry than fiction that people haven't had time for their ideas to calcify.

On the other hand, there is a definite strain in the hypertext community of folks who like to get up on a soapbox and propound all kinds of confining rules for how hypertext "ought" to be done -- I call this "fusbudgeteering". So we'll see. But I remain optimistic!

>> if you look at hypertext in terms of complete building blocks of information put together in various ways - which is how I now look at it (but perhaps you see this differently?) than it is hypertext - spatial hypertext. Are you comfortable looking at your work in this way/entirely in this way? Does it matter to you what it is called? <<

I define hypertext as any form of text with embedded interactive structure operations. But I take a pretty broad view of what those operations might be. I'm not too thrilled with the "complete building blocks" idea. It's true that that's how a lot of hypertext has been built, but I'm much more interested in work that is built from scratch as hypertext. One of the concepts I explore in my HT'96 paper

is what I call the hypertext *episode*. An episode is whatever group of activities cohere in the reader's mind as a unit. One of the ideas I'm proposing is to consider the episode as a virtual document, rather than what is known as "the lexia." For those not familiar with hypertext literary lingo, 'lexia' is a term George Landow borrowed from Barthes to describe the chunk of text at either end of a hypertext link. It is truly appalling to see how lexia-centric hypertext still is; in my paper I argue for meaning happening *through* activation of the interactive devices -- not just meaning in the lexia, but meaning in the episode. It will be interesting to see what the reaction to this concept is; so far I've gotten some pretty nice advance feedback.

My ultimate interest here is hypertext as a medium of thought. To me that *doesn't* just mean hypertext as a medium for organizing *thoughts* -- linear thoughts -- but rather as a medium in which one "thinks native" thoughts that are hypertext all the way through: hypertext extended into the fine structure of language. This is a very difficult idea, on which so far I've got very few takers.

Does it matter to me what my work is called? Not really. If someone wants to throw me out of the club of hypertext because I don't use links, so be it, but I'd rather focus on the commonalities between some of the non-link structuring methods and the more "traditional" ones. One of the things I tried to do in my HT'96 paper was to generalize some of the kinds of rhetoric Landow likes to apply to other methods, such as spatial ones. My impression is that hypertext people are increasingly receptive to this.

Jim Rosenberg

Judy:

>> And to throw in another loaded question, can you talk about language poetry and your relationship to the language poets? <<

It certainly is loaded. I hope no one minds, but since this exact issue came up on the ht_lit mailing list some months back, I thought I would post here what I said there about it, since it pretty well says what I have to say on this issue. This was posted to ht_lit (Hypertext Literature mailing list) back in May.

Michael Joyce:

>> Jim has long been associated with the L=A=N=G=U=A=G=E Poets <<

There is nothing in the whole realm of poetry that gives me more anguish than this question of what is my relationship to these people. The statement above is true in the literal sense that I've known many of the L=po folk since before there was an L=mag; we *do* share some concerns, but there are also some points of major difference. I have said it before and will say it again, even though it now sounds somewhat silly (I am not so naive as not to recognize that 'language poetry' as a term is an accomplished fact of many years' standing): the term 'language poetry' is odious, is hurtful; it hijacks the word language, which belongs to *all* poets. Poetry by those not of the L=po persuasion is made out of *what*?!

That having been said: I like the idea of thinking of language as one's material as visual artists might think of their material: something that can be manipulated, rearranged, put through *process* perhaps. I *am* interested in "language itself", I do believe my work raises questions about language, but those are *not* what the work is "about".

>> which indeed shake up the language and to realign meaning radically.. <<

As usual, Michael puts it beautifully -- certainly I would drink to this.

Where the L=po group and I part company -- drastically -- is on the issue of poetics. The L=mag established an atmosphere of poetry and poetics in close confines, on the same ground, in the same idiom even. While I can stand back and look at Charles Bernstein's formulation "Poetics is simply a continuation of poetry by other means" and whistle to myself "slick move!" I find myself taking violent objection. If you say poetics is simply a *continuation* of poetry by other means, there is a natural implication that if you don't do poetics you have cut something off, have stopped a natural continuation. This is wrong, wrong to the point of potential harm. Perhaps for Charles Bernstein ceasing to do poetics would feel like the cutting off of a natural flow that comes directly out of the work, but what is harmful here is the presumption that it applies to everyone: there is a rank prejudice rolling around in the L=po world which says: "thou shalt DO poetics -- lest thou not be taken seriously as a poet." This is just plain rubbish.

There are those who are conscientious objectors on the field of poetics, who feel -- passionately -- that the work must speak for itself, who simply don't do poetics. While I do do poetics from time to time -- in measured doses, hopefully

-- we must absolutely respect the refusal to do poetics as just as valid: it is quite literally true that in the end it is the poetry that matters; if the poetry works without the aid of poetics, there is nothing "missing" by its absence, no fault attaching to the poet, no excised connection to be looking for or wished for, and absolutely no lack of seriousness.

There is an analogy from biology that I find useful as a metaphor for the artistic process that may help explain how I feel about poetics. We all are so used to navigating by senses that operate via direct lines that we can't even imagine what it would be like to function in a world where the primary sense is chemical. Some organisms that orient to chemical gradients use a mechanism called klinokinesis. (I hope I'm getting this right ...) These animals are constantly *turning*, and the rate of turning is adjusted based on whether the turn took them to more favorable or less favorable conditions. It works well for orienting to chemical gradients, and even though these animals can't "see" the source of their food, by the laws of probability the turning is controlled in such a way that they end up at the points of maximum concentration of the chemical to which they're orienting.

This says a lot to me about the artistic process. One *is* constantly turning, and the important thing is that fine-grain sense of feedback: this worked, that didn't, this feels right, that doesn't. One can't always directly *see* hulking above the landscape the exact artistic endpoint, but ultimately the feedback from turnings enables that journey that somehow gets you to the point of maximum energy. Poetics, on the other hand, runs the risk of erecting a line-of-sight structure that sticks out above the landscape. There's nothing wrong with that, of course, but woe unto you if in the process you manage to blast those sensory nerve hairs that make the klinokinesis work. (And shame unto you if you teach young people that the klinokinesis process doesn't matter ...)

Michael writes so eloquently and often about the *topographies* of writing: having gone a journey by klinokinesis, one *may* be able to see what the topography *was* and explain it in a way that is useful to others; I suppose I tend to think of poetics as being a more retrospective activity than a prospective one. The risk is that in spite of one's intentions, the theory ends up speaking with a louder voice than the klinokinesis. Sometimes you just have to tell the theory to shut up. The non-linear work that I did prior to _Intergrams_ was a series of diagram poems going back a long way. In between _Diagrams Series 4_ and _Intergrams_ I did a whole series of good ole fashioned linear poems. I didn't really "decide"

to do that, and can't really justify it in any theoretical way -- it just happened. I did a couple of linear pieces and just kept coming back and coming back because something said that was a turn I had to make. The piece I'm working on now will probably *not* use my diagram notation. Again, this just "happened", it might go against some of the neat theory I've got -- but when the microturn says "yes, this way" you have to listen.

I suppose what I would like to keep from L=po is the inquisitiveness and energy about language as material, as process, but I don't have much affection for the rhetoric.

Judy Malloy

Thanks Jim - I'm looking forward to reading this offline. AS an aside, I was over at Poets & Writers this afternoon and I got the url for the Electric Poetry Center at SUNY Buffalo. It's http://wings.buffalo.edu/epc but I haven't looked at it yet.

Judy Malloy

>I define hypertext as any form of text with embedded interactive
>structure operations.

hmmm- this is a good definition that fits both your work and work with explicit links (Joyce, Moulthrop, Guyer) , but in that it defines hypertext in terms of navigation and structure instead of in terms of writing (or perhaps Bolter's "writing with places"), it is also a somewhat limited definition.

I would argue that using complete (screensized) building blocks of text (lexias) written with links that are implicit but not necessarily defined to the reader is also hypertext. With Forward Anywhere (the work that Cathy Marshall and I did) we were thinking in hypertext terms as we wrote in that we mentally choose something in the other's preceding text to link to. Quoting Cathy from our paper for the Wired Women book:

"But during our second meeting, we decided on a course that avoided unnecessary invention -- to exchange the real remembered substance of our lives. It is odd how the links arose quite naturally all along within this constraint, although I stewed at great length over the first screen I sent Judy -- too glib, too vague, I thought, to be evocative. The "lived lives" constraint turned out to be a positive force in the process, shaping the work in an organic way.

We agreed, at our meeting in Berkeley, to define structure later in the project, after we had amassed content by email. The links were to be left more or less implicit in our exchange (although, in practice associations often found their way into the Subject line).

In retrospect, email seems like a naturally hypertextual form, with its splitting and merging threads of conversation, its subjects that recur and re-emerge, and its tendency to discourage linearity....."

(Malloy and Marshall, Closure was never a goal of this piece.)

When we structured the work, we choose not to embed the links within the text and were somewhat uncomfortable with the semi-explicit "forward", portion of our interface (I had used this method to structure "Scibe") because the work *itself* seemed to us to be so clearly hypertext (interactive perhaps in its collaborative nature) in the process in which it was written. Quoting Cathy again from the same source:

"Besides adding the gathering function described earlier, we also decided to include a random function that brings a reader to a screen selected at random from our entire collection. The random function addresses the high interconnectivity just as surely as using a large number of explicit links, since the effect is in some ways quite similar -- a reader can get to any screen from any screen. It is the forward and backward functions -- and other explicit links -- that have given us pause. It is perhaps the sum of the experiences of past screens, the cumulative mystery, which has lead us to write the next. "

Turning this back to your work, I'm interested in how the Intergram structures evolved. How do your writing processes/your content relate to your structures?

Judy Malloy

Oh and thanks for the excellent response to the language poetry question.

Jim Rosenberg

Judy:

>> hmmm- this is a good definition that fits both your work and work with explicit links (Joyce, Moulthrop, Guyer) , but in that it defines hypertext

in terms of navigation and structure instead of in terms of writing (or perhaps Bolter's "writing with places"), it is also a somewhat limited definition. <<

Before getting into the substance of Judy's argument, there are some things I should say. I'm a pluralist, and get quite upset with people who think there is only One Way. On a certain other mailing list (that I'm not supposed to mention in cyberpublic) with lots of poets participating, they seemed pleasantly surprised when I openly stated my support for non-computer methods (like physical transparencies, etc.) I'm on record as saying that I don't think we hypertext folk are doing ourselves any good by what I describe as unguided missiles talking about "death of the book", "the late age of print", etc. In doing what I do, the last thing I want is for anyone to feel that I'm trying to *preclude* some other way of working. As I like to say, there's one house of poetry with room enough for everybody.

That having been said:

>> I would argue that using complete (screensized) building blocks of text (lexias) written with links that are implicit but not necessarily defined to the reader is also hypertext. <<

I don't know whether I buy this. Many people have argued that there are "implicit links" to a lot of different kinds of print literature, also. It's a bit painful: as with any definition, we either have to stop using the term, and admit it's useless, or we have to accept that it does place a boundary somewhere. I suppose I wouldn't mind having to give up the word 'hypertext' -- though when I'm described as a hypertext poet I don't disown that. Cage taught us to think in very large terms: The domain of music is anything which can be heard, the domain of visual art is anything that can be seen. And of course, there are times when even these terms seem too confining. To me, the domain of poetry is all of language. So yes, certainly, I'm interested in an inclusiveness that embraces "ready-made" lexia with implicit links. On the other hand, I'm not sure whether I'd call it hypertext. Maybe we need a better term. Cathy Marshall has told me that there are some well known hypertext folk who give her the impression all the time that deep down they don't consider what she does as hypertext; in that case I'd be outside the fold too.

At bottom, as a "card-carrying experimentalist", I put the questions first: *What happens if* -- what happens if I deeply interlace explicit structure operations into the fine structure of language. On the other hand, I support completely the idea of starting with the words and having the structure "migrate its way" into the words: that's the way I've been doing my work for quite a while now, actually. Everyone has to follow their own path to how the energy works; each experimentalist has to formulate her own questions. Implicit links are interesting! Maybe we need a new term that embraces everything from hypertext narrowly construed as the [notorious] node-link model, to print text that invites non-linear exploration by implicit associations. What we're really talking about here is a kind of "open-text".

Jim Rosenberg

Judy:

>> Turning this back to your work, I'm interested in how the Intergram structures evolved. How do your writing processes/your content relate to your structures? <<

I've developed what I call my "semantic method" over many many years; in some ways it works independently of the structures. I use what composers call "precomposition". That is, the piece is composed in layers of activity, and each but the final layer affects the entirety of the finished piece. I maintain what I call "reservoirs" (Jackson Mac Low calls them "vocabularies") which are precompositional groupings of phrases. At each step I take the existing generation of reservoir, chop it up, permute it using chance operations, and use the resulting "prompt sheet" to write -- by hand, so to speak -- the next-layer reservoir. At the end of this process I am writing finished phrases from the last-layer reservoir. I will typically do quite a number of pieces from a given reservoir.

Somewhere about the middle of Intergrams I began doing an "edit phase" that has become very important to me. When I have the first draft of finished phrases, I cull these. A few -- I would guess typically no more than about 10% -- are good enough to go into the finished work mostly intact, with just minor editing. The rest go into a "metamorphosis soup". (When those insects that undergo complete metamorphosis enter the pupal phase, their bodies literally dissolve into a soup of cells. Many of these cells migrate to completely different places in what will be the adult insect, and a whole new kind of organism is assembled

from the cells. That's not unlike how my edit phase works.) I pick out the pieces that still work, and then pull in words or pieces from "sacrificed" phrases reassembling new phrases. I do "cheat" occasionally -- putting in a word that didn't come from the first finished phrase draft -- but this is pretty rare.

From the result of all this, I pick the final phrases for the finished work. My wastage percentage here has gone down quite a bit; it used to be that I would throw out about 2/3 of the final phrases, but in the work I'm doing now, _The Barrier Frames_, the percentage of kept phrases is pretty high.

At this stage I move the phrases about letting them attract one another into clusters, with the structure emerging from the words. Composition of the words has normally taken about 85% of the elapsed time of making the work.

With _The Barrier Frames_ things have worked out somewhat differently. I composed simultaneity structures in the small as I went, then redid them using the same edit phase mechanism; in the large the clusters emerged from the finished words, as before. Also in this piece, the "assembly" phase has been taking me *months* -- it's excruciating!

Some of these methods are compromises based on the available tools. I simply don't have an authoring system that will let me make reservoirs out of pieces with the interactive structure operations built in, which I can then "plug in" elsewhere. I have dreams of an authoring system that will let me do this, but don't have time to program this myself. (When I was an undergraduate, I knew a guitar maker who lived in the town where the college was; I would hang out there occasionally. I was naive enough back then that when he said he couldn't afford his own guitars, I was shocked. I sure understand it now: I earn my living as a programmer, but as far as using my programming skills for large projects involving my writing, I can't afford myself!)

Judy Malloy

>Everyone has to follow their own path to how the energy works; each
>experimentalist has to formulate her own questions.

Yes! This is by far the most important thing and I agree 100% Every writer, every artist has his or her own way. That's one of the reasons print is not going to disappear. This week Michael Joyce's phrase "print stays itself; electronic text replaces itself." came to mind. I realized that it is also an apt definition of your work, (Do you think so?) but it is nothing like my way of thinking.

> Maybe we need a new term that embraces everything from hypertext
> narrowly construed as the [notorious] node-link model, to print
> text that invites non-linear exploration by implicit
>associations. What we're really talking about here is a kind of
> "opentext".

For me the ability to manipulate the text is very important. That manipulation is what computers make possible. I'm less interested in so-called print hypertext - perhaps because I spent some many years trying to make nonsequential text work in print so I know first hand that however you structure it, it will be read sequentially. I'm rereading Hopscotch (Anna gave it to me for Christmas!) and I don't skip around in it. I read it from cover to cover. It is an excellent book, but I feel no kin with its structure.

>On the other hand, I support completely the idea of
>starting with the words and having the structure "migrate its
>way" into the words: that's the way I've been doing my work for
>quite a while now, actually.

Working on the Yellow Bowl has been an epic interface battle - symbolic perhaps of the difficulties of attempting to combine sequential and nonsequential text and of the primary importance of words and intention. YB was designed (I also use a form of precomposition that involves large charts on my walls) with what I conceived of as three streams of text - two were sequential (the stories that the narrator told her child) and one was pseudorandom (the narrators thoughts/ memories that she distorted to shape these stories -designed to appear in changing but natural memory patterns) The reader was supposed to move in and out of the stories and the narrator's memory bank. I used peripheral "frozen" links. ("buttons" - I hate that word)

In installation situations, I discovered that people didn't bother to get into Grace's mind. They stayed in the stories til they got to the end. So, I experimented with arbitrarily throwing them into the memory bank at certain intervals and tried it again. They got out as soon as they figured out how- the human desire perhaps to know the end of the story. But in this work, the stories by themself

were not important. It was how they were shaped from the narrators experience (and how parents communicate experience to their children) that was backbone of this work.

So had to change the interface. (although the underlying structure remains the same) The work was already written to links (implicit to the reader but explicitly written to) so it was easy to highlight them using a "linkplot" strategy with dense linking between the three "streams" of text that guided the reader with links that sometimes produce the random material and sometimes continue the story. Currently, YB is more "authorial" than I would like, but I think that it now "works" in a fairly seamless way. However, I wouldn't make changes because of how readers experienced the work in every one of my works. In this work, the reading pattern was central to my aims.

Poetry is different than narrative. The roots of your work, the writing of your work are very different from the roots and the writing of my work. Nevertheless, I'd be interested in how you have observed readers interacting with the Diagrams. Has reader experience has shaped your work in any way? Is that is important to you?

>what happens if I deeply interlace explicit structure operations into
>the fine structure of language.

I would like to get more deeply into what you mean by this. Feel free to plunge in if you want. Meanwhile I am hoping some of the musicians on Arts Wire will to address your endlessly fascinating composition methods.

Judy Malloy
oh and I meant interacting with the Intergrams.
Jim Rosenberg
Judy:
>> Poetry is different than narrative. <<

Right. I don't have "the" story to which I have to hold any allegiance.

I suppose, as a reader, I'm something of an oddball. I think to most people the essence of a good book is one that you can't put down, you have to keep reading to find out what happens. To me, the essence of a good book is that you *have to* put it down because the imagery is so rich, is resonating so fully, that you have

to put it aside and let it ring for a while. Resonance is clearly a whole different animal than narration.

>> Nevertheless, I'd be interested in how you have observed readers interacting with the Intergrams. Has reader experience has shaped your work in any way? Is that is important to you? << [Correction Diagrams -> Intergrams interposed ...]

I suppose perhaps my answer here will bother some people, but no, I haven't "studied" what readers do with _Intergrams_ or _Diffractions through_. I get a curious feeling when I'm present and someone is operating my work; I feel as though I'm in danger of eavesdropping on what should be a private moment between the reader and the work, and always have a *very* strong inclination to want to leave, to leave them alone with the words. I can see where some people would feel that part of the job is studying the ergonomics of one's own creations, but I see it differently. When you get past all of the technology and the buzzwords, what it's all about is nothing more complicated than simply:

bringing words together

and ultimately, as an artist, I have to rely on my own convictions and intuitions of what's right and what's not right. I like your idea of studying what readers do to try to "throw them out of the narration channel" -- maybe you wouldn't like how I've put that -- but I've never felt impelled to do anything along these lines.

Of course as I'm making new work, there is a handful of people whose opinion matters a lot to me; I will bug them about whether things "work" or not, and may have an occasional precise question or two about whether specific interface features work the way I had in mind, but I don't really "study" what readers do with my work.

I'm really eager to see what you're doing in the Yellow Bowl.

Jim Rosenberg

>>> what happens if I deeply interlace explicit structure operations into the fine structure of language.

>>I would like to get more deeply into what you mean by this. Feel free to plunge in if you want. <<

This is a large subject, which I've been grappling with for a very long time and about which I care very deeply. The rudiments of my diagram notation first came to me somewhere along about 1968; at first I simply conceived of it as a device for making poetry, but in the summer of '69 I had a flash one night at about 2 in the morning and realized that "language itself" could be based on these concepts. In all the years since I've been waging an ongoing battle with myself to try to get the ideas more and more explainable, but it is a really difficult concept. Eduardo Kac is guest-editing an issue of _Visible Language_ that will have papers in it by himself, the British computer poet John Cayley -- whose work I admire very much and several others; I have a paper coming out in that issue called "The Interactive Diagram Sentence: Hypertext as a Medium of Thought" that goes as far as I've been able to get lately in explaining this. I wish I could post that entire article, but I don't know whether I can do that. Most people approach hypertext from what I would call the outside in. That is, they start with whole documents -- typically documents written "pre hypertext", with no interactive structure at all. They then begin to impose more and more interactive structure on the document; perhaps section navigation, table of contents, etc. are handled through links, a few key words or topics are made into link anchors, then gradually more and more links are made, and so on. The focus is still "the document".

By having started with my Diagram Poems, I was going about it exactly in the opposite way. We normally think of hypertext as software, or as the electronic documents created using software, but I think of hypertext as a virtual *diagram*. You certainly get a diagram in a conventional node-link hypertext if you diagram the link network. The Diagram Poems were a sort of proto hypertext "from the inside out". The diagram structure forms the very basis of the syntax; rather than starting with linear documents and using "hyper" operations to connect whole documents, or sections of documents, the very sentence itself is constructed using these methods.

It's very important in understanding how this works to keep in mind the and/or distinction, or what I've called conjunctive vs. the more typical disjunctive hypertext. In the usual approach to hypertext, the reader has a choice among *alternatives* in choosing which link to follow. If lexia X has links A, B, C, the reader can choose A *or* B *or* C. I've been arguing that it also makes sense to build hypertext based on "and" rather than "or". In logic, "and" is a conjunction, "or" is a disjunction, so I call this conjunctive hypertext rather than disjunctive hyper-

text. (The paper where I introduced all of this is Navigating Nowhere / Hypertext Infrawhere) Conventional syntax is conjunctive: if we say a sentence is composed of a noun phrase and a verb phrase, the noun phrase and verb phrase are hardly *alternatives* -- neither one is "optional", they're both there, you have to have the noun phrase *and* the verb phrase to get the sentence.

Sentences are also a kind of virtual diagram. You could say that syntax has been nature's way of allowing us to fit very complicated structure relationships into a word stream that is constrained by two important limitations: (1) Throughout most of the time speech has been existence, speaker and listener have had to be synchronized in approximately real time. (2) The word stream itself has no storage; storage has to be achieved by speaker and listener. You could say that syntax is a way of encoding *storage cues* into the message so that structural relationships among distant parts of the word stream can be constructed in the face of these constraints. But writing does not have either of these constraints. There is no time synchronization constraint at all with a written document, and the document contains its own storage. In writing we are free to indicate structural relationships *directly*.

This is what I began doing in my diagrams. Of course the problem with a diagram is how do you get it to scale into an immense number of elements. This is where hypertext comes into the picture. Hypertext lets you build a virtual diagram of unlimited size. These hypertextual operations can carry the same structural relationships as syntax -- *inside* the sentence Meaning is not just a function of the lexia, but happens has we move *through* the links, at the level of what I call the *episode* in a paper I'm giving at the Hypertext '96 conference in March.

The issue here really is hypertext as a medium of thought. By this I don't mean hypertext as a medium for organizing *thoughts* -- linear thoughts which are not themselves hypertexts, but rather as a medium for "thinking native" in hypertext, where the *individual thought* is a hypertext, "all the way down". In the _Visible Language_ article I raised the question of what would it mean to construct a *multiuser sentence*. That raises the issue of hypertext inside the sentence in a context where I don't see how to do it any other way.

How do we construct a medium where *an individual thought* can be a multiuser construct?

Judy Malloy

Thanks Jim! I'm really enjoying this conversation and am looking forward to reading and responding offline. Thanks again for taking the time to visit us.

Anna Couey

Jim, there's a lot you've said all thru this topic that has sparked much thinking. I suppose by your definition of a good book, this is a good topic, because I want to absorb it a while before tossing a response back at you.

About the levels of interactivity in Diagrams & Intergrams. Your answer seemed so obvious when you said it. It wasn't until later that I remembered what may now be an art expression rendered technologically obsolete - that the choosing of an online work from a menu was in effect causing it to appear. And we considered it more interactive than a static physical art work because the work wasn't there until the viewer called it into being. Intergrams carries reader choice into the poems. It is several steps deeper in, and is certainly more flexible than the interactivity of Diagrams. The reader is not physically altering your work - that is, each reader will see the same work, albeit through different paths. Is the interactivity in the ability for the reader to select a path through the work?

A short question about your last post...have you explored whether hypertext as a medium of thought exists in other languages?

Judy Malloy

>>To me, the essence of a good book is that you *have to* put it down because the imagery is so rich, is resonating so fully, that you have to put it aside and let it ring for a while. Resonance is clearly a whole different animal than narration.<<

Resonance is a vary good word that is applicable to hyperfiction as well as poetry. I think in terms of small very rich pieces that the reader assembles in different ways but that stand on their own because of what you call their "resonance".

assembles I like the way you put this:

<<<

>>I've been arguing that it also makes sense to build hypertext based on >"and" rather than "or".

I tried this using a frozen interface with selected keywords in the original second file of Uncle Roger because coming from working with databases I was accustomed to using "and" as well as "or". In this work (The Blue Notebook) the reader could make a path by combining several words. The problem was that reader "searches" often produced minimum results or no results. In a narrative, this interrupts the flow. (Although I can now see instances where this interruption could be a part of the narrative strategy, and in addition I am now more experienced in writing in tandem with the structure) Actually I think it would work much better with poetry and in particular I would very much like to see how you would apply this to your work.

>you have to have the noun phrase *and* the verb phrase to get the
>sentence.

And I never thought of it in this language structure way. This sounds very interesting.

>>Meaning is not just a function of the lexia, but happens has we move *through* the links, at the level of what I call the *episode* in a paper I'm giving at the Hypertext '96 conference in March. <<

Yeah - I tried conventional linking with l0ve0ne and found that in the on the click-happy web, it was being read in this way - click click click. It was a chain of linked phrases that stuck in the reader's mind and to a certain extent I began to take this into account as I wrote.

>>The issue here really is hypertext as a medium of thought. By this I don't mean hypertext as a medium for organizing *thoughts* -- linear thoughts which are not themselves hypertexts, but rather as a medium for "thinking native" in hypertext, where the *individual thought* is a hypertext, "all the way down". In the _Visible Language_ article I raised the question of what would it mean to construct a *multiuser sentence*. That raises the issue of hypertext inside the sentence in a context where I don't see how to do it any other way. <<

I have been in therapy because of post traumatic stress syndrome and have had to confront how my own thinking about life/thought processes have

merged with hypertextual writing. This is slightly different from what you are talking about, and I'm not quite ready to put it in words. Nevertheless, it is very interesting.

>How do we construct a medium where *an individual thought* can be
>a multiuser construct?

Can you talk a little more about this?
Jim Rosenberg
Anna:
>> ... that the choosing of an online work from a menu was in effect causing it to appear. And we considered it more interactive than a static physical art work because the work wasn't there until the viewer called it into being. <<

Oh dear. I didn't think of what we were doing on ACEN in these terms *at all*! It simply isn't true that the work "wasn't there" until the viewer "called it into being". Maybe it's my programming background -- the fact that I know what's really happening in the system -- but I simply can't think of it that way. All those menu choices that the reader hasn't made *are* there: they are files in the system. The reader's choice doesn't call them into being, menu choice is just a more convenient method of saying "cat /home/j/e/jer/diags/17", but the user doesn't "create" that file. The user asks to have the file *transmitted*. In a way, it's no different than special-ordering a book and asking to have the book *transported*. Of course the latter is a manufactured object, but in the former case there is no manufacture unless the user prints it on her own machine.

>> The reader is not physically altering your work - that is, each reader will see the same work, albeit through different paths. <<

Yes: this is the distinction Michael Joyce calls Constructive vs. Exploratory hypertext. There are those in the hypertext community who really push for constructive work, where what the reader does actually changes the work that *others* will see. The work is a genuine collaboration between readers and writer; in effect all are writers. There is a lot of hypertext rhetoric about reader-as-writer, to

the extreme that in some cases one gets the idea that works that are only readable without the reader being able to change them are somehow "inferior".

This sentiment causes me a lot of anguish. I like the idea of constructive hypertext; comparable to what Judy was saying about the non-linear aspects of all writing, I think of meaning as a constructive act, even if the reader is not literally altering the work, and am pleased to see experimenters pushing the envelope by making works where the reader is a participant in every sense of the word. From time to time I think that I should be working this way too, but when I ask myself, "OK, *what*?" -- what would I like to do -- I don't come up with anything. In some ways I guess I'm a loner at heart. There are many people these days who push very hard a collaboration ethic that seems to imply that collaborations are inherently "better" than individual works; I think we need *both*. Collaborations haven't really happened much in my artistic life, it's just how things have worked out.

>> Is the interactivity in the ability for the reader to select a path through the work? <<

Sort of -- though I'd step back a bit from the word 'path', since that implies a choice among alternatives (or vs. and again ...) I think of the interactivity this way: I've presented the reader with word objects, where intelligibility can only be recovered by *doing things*, by "operating" the objects.

>> A short question about your last post...have you explored whether hypertext as a medium of thought exists in other languages? <<

That's a really interesting question. I've often thought that a language like Chinese would be better suited to non-linear writing than English; because the word is so spatially compact, it provides more room to maneuver, spatially. It turns out there is a tradition of non-linear poetry in China that goes back 1500 years!! There is a work known as a Poem Block, attributed to Su Hui from the 4th century A.D., where the characters are arranged in a rectangle that can be read in practically any way; *hundreds* of such readings have been constructed from her block. I think of this from time to time when I'm in danger getting too uppity

about how wild and new what we're doing is. It's pretty humbling to think the Chinese were doing this 1500 years ago!

Jim Rosenberg

Judy:

>> I think in terms of small very rich pieces that the reader assembles in different ways but that stand on their own because of what you call their "resonance". <<

Yes, absolutely: assembly is a *wonderful* word!

>> > How do we construct a medium where *an individual thought* can be a multiuser construct? <

Can you talk a little more about this? <<

I guess it's a bit ironic, given what I said in the previous message about collaboration ...

We tend to think of the sentence as a "single voice". Even when a work has multiple authors, the usual method is to strive for a "collective voice", so that each sentence sounds as though it has a single (albeit collective) author. In a diagram syntax, a single sentence can be assembled by several people; the contribution of different people could be indicated say by color. I suppose in a way this isn't a lot different than a hypertext in which different links and anchors might have different -- and indicated -- authors, though it's carried out at a very fine scale, inside the sentence.

Almost everyone these days seems to understand how huge the influence of networks is on our lives: they are everywhere, they pervade all facets of our lives. Let me pose a very naive- sounding question: How does a single mind apprehend a network? Might it not help to *become* the network? Might it not help to "think native" in a network language, in a language in which an individual thought *is* a network?

This is really at the heart of what I'm talking about when I speak about taking hypertext into the fine structure of language.

Judy Malloy

Thanks, Jim. It has been wonderful having you here. Is there anything that you particularly want to talk about?

Jim Rosenberg

I've enjoyed this immensely! There simply is no way to sum up the issues we've been talking about; one simply has to be clear that it's the questions that matter, that this is drastically unfinished business, and hope everyone will keep on experimenting away. When I first joined ht_lit -- and one of the first questions here! -- the issue was defining hypertext. Actually, my favorite answer to that *really* is: "We're working on it ..."

That seems like a pretty good note to end on. Thanks to you, Judy and Anna for a wonderful discussion.

Transcript of A Conversation with Jim Rosenberg, Item 74, Interactive Art Conference, Arts Wire.

Posted to the Web with participants' permission.

Conversations with Artists

INTER-FIELD, THE ACTS

[The Field Project is a collaborative essay by mIEKAL aND, Jim Andrews, Thomas Bell, Loss Pequeño Glazier, Inna Kouper, Clemente Padin, Jim Rosenberg, and Ted Warnell.]

Field. As in Composition By Field -- a term I have always attributed to Robert Duncan, from whom I heard it first. I have always taken Composition By Field to mean this: in the poem a locality or area or field becomes established as the poem is being composed; if a continuation occurs to the poet which works within the field, then it should be accepted as a valid continuation of the poem without having to resolve it against any kind of global constraint or global set of rules. Thus, Composition By Field is a kind of composition by induction (in the sense of mathematical induction). Mathematical induction is a method of proof by continuation: you show that a statement is true initially, then you show that if the assertion holds, it follows that it continues to hold for the next number. This establishes that it holds for all numbers.

Continuation: producing the nextness. What does this mean for a poem that is structured nonlinearly? Surely in hypertext we have a strong sense of locality; the term lexia has become entrenched in the hypertext vocabulary: the location where you are when you go nowhere. (Which is not to be complacent in the face of arguments that 'lexia' is a concept which is seriously problematical -- see "Navigating Nowhere" below for a discussion of circumstances in which the concept of lexia may be said to break down completely.) Surely continuation is something that can happen through the link, not just within the lexia; link-based hypertext has a strong concept of continuation, except that it becomes pluralized: there may be a multiplicity of continuations.

Along with composition by induction, the poetics of the 60s had a strong sense of Real-Time Poetics: the time sequence of events in the poem should have a direct relationship to the time sequence of events in the composition of the poem. Read-time continuation had a relationship to write-time continuation. Allen Ginsberg wrote about this explicitly in Improvised Poetics. Here we begin to get into a very serious problem with nonlinearly structured work. In a link-based

hypertext, read-time continuation may go through a given link. Write-time continuation may have gone through a different link. Even when simply focusing on the reader, the issue of continuation becomes complex. There is a multiplicity of continuations. For the reader to experience this multiplicity, continuation must be a separate concept from contiguity in time. To experience the multiplicity the reader must return, perform a different continuation. Perhaps the time of these continuations is equivalenced. (Though one must admit this concept of equivalenced time is an abstraction.)

This is not to say that multiplicities in hypertext should always be taken as alternatives. The elements of a multiplicity may be taken as components. A unit may be a cluster consisting of layer A and layer B and layer C as opposed to a lexia from which one may take link A or link B or link C. Conjunctive hypertext. (See "Navigating Nowhere" below.) (In logic a conjunction is and, a disjunction is or.) For the linear field as practiced by Robert Duncan, the continuation happened in a direction, onward in the poem. In a conjunctive cluster directionality is not onward but inward, toward an abstraction of the cluster as something in a fixed spot, a center, a locality. The field is a real field: an arena. An arena in which acts occur on the text, perhaps in which the text itself acts. Word objects that behave.

How do fields combine? For the linear poem composed by field, the reader simply continues reading -- or the writer simply continues writing. It is not a coincidence that the poetics of the 60s made a strong connection between fields and time sequences. How do you combine fields whose directionality is inward? Fields can be combined by just doing it: by spatial juxtaposition, by drawing the connection explicitly, by embedding a device on the screen area where the combination happens through programmed behavior. Space replaces time. The field is a container. Fields can be combined by nesting them, or stacking them, or adjoining them, or doing the virtual equivalent of connecting them with wires. Fields connect the way parts of a sentence connect: with plugs and sockets.

Fields combine in a network. (The linear chain may be thought of as a special case of network -- mathematicians would call this a "degenerate case.") The network may have "disagreeable" properties. For instance it may contain loops -- feedback loops perhaps. The feedback loop is a difficult structure. Program-

mers typically cannot deal with such structures: they like their structures to be "well founded," parts having other parts etc. in a clear hierarchy, so there is a clear bottom level nesting back to the top. Computer programs normally deal very poorly with feedback loops. But the mind can apprehend a feedback loop "at once" -- as a single gestalt. Surely the feedback loop is one of the most fundamental structures of nature. The idea of a poem which folds back on itself at the end is a venerable one, one of the classic ways to end a poem. In the network way of thought, a thought itself is a network, down to a fine grain. The network of fields builds field combinations without reservation. Proximity need not matter. Topology need not matter.

Replacement of one field by another. (Michael Joyce's concept of writing that rewrites itself.[1] Continuation not by the movement of the reader "in" the text but by the text itself in front of the reader. The lexia is replaced by a new one. There are some difficulties here. The replacement is typically complete: no residue of the former lexia remains in view. No one ever said that fields in the classical Composition By Field couldn't overlap. In fact you could almost count on the fact that fields would always overlap: the word is always at the center of a field: energy radiance. Perhaps we need new interface ideas in hypertext, a moving-frame lexia.

Reading through the network needs to be not so much a navigation as a gathering: bringing back to central hinge points the results of the fields, integration of the multiplicity. The technology is still in its infancy at providing good gathering tools: let no one think that at this moment either the interfaces or structures for hypertext are a done deal! We are all learners at this. (We are all also latecomers: Lady Su Hui produced in The Revolving Chart [Fran] in China in the 4th Century A.D. an intensity of connectivity we have yet to recover: a matrix of ideograms that can be read in almost any direction -- complete connectivity -- with literally hundreds of paths known in the Chinese critical literature.) The field is quite literally a Field: the place of harvest. The locus of gathering.

1 See Joyce's *Of Two Minds: Hypertext Pedagogy and Poetics.*

Inter-field, the acts: the field in one's fingers, mouth, lips, jaws, the whole body. The acts: the doing: it is an act that replaces field with field, an act as part of a multiplicity. Interact. The poetics of the 60s had an anatomical fixation on breath, as if there are no words outside of breathing. Multiplicity of breath in a single body is a difficulty. The eyes make multiplicity easy, the movements of eyes are outside of our awareness. Multiplicity of events in the world simply happens, we do not breathe it. There is too much unity in a single person's breath: the breath is not a world. Inter-field: the fields so multiple, so dense, the poem is a world. A word forest. Field to field is a multiplicity of breaths: not the breath of the poet, the many breaths of the world. An ecology of breaths.

A PROSODY OF SPACE/NON-LINEAR TIME

PART I: BACKGROUND: LINEAR PROSODY[2]

DIMENSIONS OF INEQUALITY AMONG SYLLABLES

Prosody in the English language proceeds from the axiom that not all syllables are created equal; many effects in prosody derive from the time-plot of these inequalities along various dimensions. The most well known of these is the familiar stress-degree, but I will quickly review others.

PITCH-DEGREES

The usual approach to pitch in prosody is to consider it a "curve": the intonation curve. However, there is a manner of recitation at work in many American communities, most notably in a style of reading in the black community, in which tight-knit patterns of time of various pitches are articulated, in much the same way that stress occurs in more traditional prosodies. This is a very rich prosody that deserves to be studied in its own right. A predominantly pitch-degree prosody will have very different characteristics than a predominantly stress-degree prosody. Pitch is a purely acoustical property, as opposed to stress, which is a linguistic property that is quite difficult to define acoustically. Thus a pitch-degree prosody is much closer to music (in the literal sense of the term); a pitch-degree prosody is freer to use an absolute musical sense of time, whereas a stress-degree prosody is more likely to be based on "linguistic time," which works differently (see footnote 8 below). Not all phonemes carry pitch; a pitch-degree prosody may thus change the sound structure balance for how phonemes relate to one another. Where both pitch degree and stress-degree prosodies occur simultaneously, incredibly subtle effects are possible.

---

2 This section is a revised version of the first part of my "Notes Toward a Non-Linear Prosody of Space" (1995). A version of this paper was presented at the Assembling Alternatives conference at the University of New Hampshire in September, 1996. My thanks to Romana Huk for that opportunity.

VOWEL POSITION DEGREES

In explaining the meaning of the term "Tone Leading Vowels" as it pertained to the prosody of Ezra Pound, Robert Duncan explained the term as meaning two things: (1) Where a diphthong (a glide between one "pure vowel" and another) occurs, the leading pure vowel of the glide plays a special role. (2) A sound is reinforced when you hear it again, but can also be reinforced when you don't hear it again. A similar concept to this second point is the idea that vowels form clusters according to the position of the mouth when they are articulated; the tight-knit pattern in time that delineates which of these clusters is active can form a prosody, much like the stress-degree or pitch-degree prosody.

STRESS-DEGREES: CLASSICAL PROSODY

The most familiar basis for metrics in English is the tight-knit pattern in time formed by stress-degrees. Stress has been extensively studied in linguistics.[3] Before introducing an alternate methodology for how metrical studies of contemporary poetry might be conducted, I will review briefly the traditional account of how the stress-degree metric is supposed to operate. This account has become a significant obstacle in pursuing prosody of contemporary poetry, so it would be well to understand it before considering a different approach. Classical prosody starts with an a priori inventory of templates of stress-degree patterns (e.g. iamb, trochee, anapest, etc.). "Scanning" is the process of matching these templates to the poem; where repeated instances of a single template match, end-to-end, the line or poem is said to "scan." It is important to note that the word "foot" is profoundly ambiguous in this process, having at least the following two meanings: (1) We speak of a foot as meaning one of the templates. In this usage, "foot" is an abstract concept which exists in advance of any particular poem. (2) We may refer to the actual syllables in a poem matched by a template as being a foot. In this usage, "foot" is a part of a living, breathing poem--and as such is a unit of rhythm intermediate between the syllable and the metric line. Much of the

3 See Chomsky and Halle's *The Sound Pattern of English*.

poetics that has been influential since the fifties and sixties has focused away from the a priori (Olson, Ginsberg)[4] and many contemporary poets are uncomfortable with the idea of a template-based metrics. Most poets and many theorists have turned away from the study of metrics, rather than explore the second usage of "foot" in which the unit of metrics is not thought to exist prior to the poem, but is rather part of the poem itself, intermediate between the syllable and the metric line.

Thus I turn now to consider this concept of an intermediate unit of meter, one that de-emphasizes the a priori and does not use any concept of template. To avoid confusion, I will abandon the use of the word "foot" and instead use the term "measure."

BONDING STRENGTH

Another dimension of inequality among syllables (really of syllable boundaries) is "bonding strength": the degree of attraction of a syllable to the one ahead of it or behind it. Bonding strength may be defined as the extent to which an artificially injected pause at a particular syllable boundary seems natural or not when compared to the way the poet would typically recite the line. Syllable boundaries will differ in their degree of bonding strength; by collecting together into a single unit those syllables where the bonding strength is high, one obtains a "measure." It cannot be emphasized strongly enough that the assessment of where the measure boundaries are located must take place with respect to a particular recitation--presumably the poet's. A printed text of the poem on the page may not give sufficient information without a sound recording. In this methodology, scanning consists of identifying where the measure boundaries are, where the rhythmic line boundaries are (a rhythmic line is a cluster of measures connected by somewhat higher bonding strength, just as a measure is a cluster of syllables connected by the highest degree of bonding strength), and then attempting to discern whether there may (or may not) be any regularity to how measures are constructed. Thus rather than speaking of a poem as being "written in" a meter, meaning a conscious a priori choice of template, one examines the poem empiri-

4 See Olson's "Projective Verse" and Ginsberg's *Improvised Poetics*.

cally to determine whether there simply happens to be some regularity to the way the measures are constructed.

THE "STANDARD MEASURE"

This methodology need not be restricted to poetry: any recitation can be scanned. The statement is often made that English is iambic. Using the method sketched above to determine measure boundaries, we can reformulate the tendency of English toward the iambic, without exempting the many counterexamples. Measure boundaries in English prose tend to be constructed as follows: (1) a measure has only one major stress; (2) the measure tends to end on a major stress, but: (3) if there are unstressed syllables following the major stress out to the end of a major grammatical unit, those unstressed syllables will also be incorporated into the measure. Measures constructed in this way may be called "standard measures." Of course not all measure boundaries in poetry will be standard measure boundaries: Robert Creeley, for instance, is well known for having many non-standard measure boundaries in his poems. Interestingly, when Creeley's poems are actually scanned, the results show that while there may be non-standard measure boundaries at the end of the rhythmic line, many lines contain two measures, and in these lines the internal measure boundary is a standard one: the celebrated Creeley line-break really is a line-break and not a measure break. The non-standard measure boundaries are very easy to hear, but the internal standard measure boundaries are much more subtle. Of course if they were missing, we would certainly hear the result as a flat, lifelessly too regular, much less interesting rhythm. The structure of Creeley's lines may be described as an "offset structure": the sound structure of the line endings is clearly articulated, but the grammatical structure proceeds from the middle of one line to the middle of the next. The offset structure is an extremely venerable structure in prosody, going back at least to Anglo-Saxon times.

PART II: NON-LINEAR PROSODY

BONDING STRENGTH IS SPATIAL

I have described bonding strength as the attraction of a syllable to the syllable ahead of it or behind it. Although prosody is normally interpreted as how the sound structure works in time, clearly the concept "adjacent" is a spatial concept; thus bonding strength may also be interpreted as a spatial concept, and as such can work in any topology, including a non-linear one. Where above I defined bonding strength as the tendency of a syllable boundary to resist injection of an artificial pause (a time concept), we could as easily have described it as the tendency of a syllable boundary to resist injection of space. It should be noted that in one dimension, space and time are nearly the same thing; however, in the more complex topologies of non-linear writing, as we shall see, space and time operate very differently.

A REVIEW OF HYPERTEXT STRUCTURE TERMINOLOGY

I have introduced a framework for structuring hypertext activity elsewhere and will review it only briefly here. By hypertext I mean a text that contains embedded interactive operations when considered from the reader's point of view: the text contains interactive devices that trigger activities. The most familiar of these is the hypertext link, but many other types are possible.[5] For instance, my work often contains devices called "simultaneities," in which groups of words are layered on top of one another; by moving the mouse among no-click hot spots, the different layers are revealed. Research hypertext software has been built based on both set models and relation models, and spatial hypertexts have been constructed using such concepts as piles and lists. In all of these cases, the hypertext is operated by performing activities; these activities consist of such actions as following a link, opening up a pile or simultaneity, etc. I have called these small-scale activities "acteme."[6] In the node-link model of hypertext, the acteme of following a link may be described as "disjunctive," from the logical term disjunction, meaning "or." A disjunctive acteme presents a reader in a given position

5 The advent of the World Wide Web has benefited hypertext immeasurably, by vastly increasing exposure of hypertext to a truly mass audience; however it is regrettable that the limited forms of hypertext activity currently available in HTML limit understanding of the variety of hypertext activities that are possible. Some of these limitations can be overcome by extensive use of richer Web languages such as JavaScript and Java.

6 See Chapter [Chapter Number],"The Structure of Hypertext Activity" below.

in the hypertext with a choice: she may follow link A or link B or link C. Other forms of acteme may be described as "conjunctive." A conjunctive acteme such as a simultaneity with layers A, B, and C consists of A and B and C.[7] A given hypertext can use both kinds of actemes together and a hypertext poem could even blur the distinction between them.

In most cases, the text in a hypertext appears in units called "lexia," a term of analysis George Landow borrows from Roland Barthes to apply to hypertext. In a typical node-link hypertext, the lexia is the unit of text at either end of a link; often (though not inevitably) the lexia has an internal structure which is simply linear. As we will see, particularly in the context of poetry, the concept of lexia is extremely problematic.

As the reader navigates a hypertext, activities will (hopefully) cohere together into units called "episodes." For a node-link hypertext, the episode will tend to be all or part of a path. It must be noted that not all activities will necessarily resolve into an episode. Some activities might be performed by accident, as when a reader pulls down a menu of link names and chooses the wrong one unintentionally. A reader may backtrack, having decided that performing an activity got nowhere. (Backtracking is complex; it may or may not revoke membership of an acteme in the episode.) Thus, episode is not the same thing as history. At a certain point the reader may not have constructed an episode at all, and might indeed be best described as foraging for an episode. The episode is an emergent concept; it emerges retroactively. Ideally, the structure of episodes emerges through the use of a "gathering interface." Unfortunately, available gathering interfaces are still quite primitive: they construct something more akin to the bookmarks of a web browser than a full picture of hypertext activity.

PROSODY WITHIN THE LEXIA

7 At its most extreme, hypertext structure may be used to represent the structure of syntax itself. In this case one clearly has conjunctive structure: a sentence consists of all of its parts; e.g. if we describe a sentence as consisting of a noun phrase and a verb phrase, the noun and verb phrases are hardly alternatives.

In many cases—perhaps most cases—the lexia is structured linearly. Under these conditions, within-lexia prosody includes traditional linear prosody. Not much need be said here; indeed one would be hard put to make the case that there is any difference between within-lexia prosody for a linearly structured lexia and the prosody of the printed page. However, there is no reason at all to suppose the lexia must be linear[8] (on the linearity of lexia see Moulthrop; Rosenberg, "Navigating"). In this section I move to consider within-lexia prosody for a non-linearly structured lexia.

Consider Figure 1, which shows a single screen from a simultaneity taken from one of my works (Rosenberg, Diffractions). This screen can be read in at least two different ways: (1) It can be read polylinearly so that the words with the same font are read as a linear skein, beginning with the word that is capitalized. (2) Alternatively, the graphically clustered fragments of these phrases can be read in snatches as the eye wanders about the surface of the screen picking up groups of words and associating them in whatever way seems to work. Even a simple polylinear reading poses difficult questions for the concept of lexia: is the lexia the entire screen, or one of the skeins? A computer-oriented view of the lexia would tend to regard the lexia as whatever is visible on the screen when there is no input to the computer, when the mouse is not moved, and no key is pressed. In this case the entire screen should count as one lexia. But what happens, in terms of prosody, as the eye moves from one phrase to another? Is this time which "doesn't count"—a kind of time out, in which there is no prosody?[9] If indeed the time between phrases doesn't count, we may describe the time units within the skeins as disengaged from one another. Or perhaps the prosody of the individual skein, together with the layout of the screen, helps determine when the next phrase begins, in which case the time between skeins definitely is part of the prosody.[10] A lexia with this type of polylinear structure is inherently ambiguous

8 On the linearity of lexia see Moulthrop's "Toward a Rhetoric of Informating Texts" and Rosenberg's "Navigating Nowhere," below.

9 Gerard Manley Hopkins defined an outrider as a syllable that "doesn't count" in the prosody. I must confess to not understanding the idea of a syllable that doesn't count. The idea of an emptiness that doesn't count is easier to understand, but that, too, seems problematic.

10 In "A Note on the Methods Used in Composing the 22 Light Poems," Jackson Mac Low instructs: "The empty spaces in 'Asymmetries' are notations for silences lasting at least as long as it would take the reader to say the words printed directly above or below them." A similar approach might leave a silence between units equal to the length of the last measure encoun-

concerning the prosody of what happens between phrases. Still another possibility is simply to say the time relationship between phrases is in the reader's hands completely. Of course something will happen when the poet recites such a lexia: a choice will in fact be made. In this case, the poet may experience a contradiction between her desire to present the work in a context where oral experience is expected and her desire to leave open as many options as possible for the reader.

Figure 1.

These issues become even more difficult if we use method 2 to read this screen. What is the prosodic relationship between these clusters of words, read by a kind of "visual wandering"? In this case linearity is so seriously fragmented that the reader may have an impression of the words disengaging from time altogether, such that prosody relationships become entirely spatial.

PROSODY THROUGH THE EPISODE

There is no reason to assume that prosody should be confined within the lexia. In this section I explore issues of prosody within the episode as a whole that go beyond the boundaries of the lexia. "Text" occurs in many places in a hypertext besides the obvious text in the lexia. There is also text in the devices of the hypertext mechanism itself. For instance, many hypertext systems allow the user to bring up a menu of possible outgoing links. Such a menu is inarguably textual. But what role does such a menu play in prosody?[11] One approach is to consider the menu of link names as a text object in its own right. Hypertext poet Deena Larsen constructs poems from assembled link names.[12] This approach, while interesting, simply reconstitutes the menu of link names as a different form of lexia, though one that has a complex structural relationship to the lexia from which

---

tered, or the last rhythmic line. A directive "leave whatever silence between units seems natural" might tend to resolve to one of these possibilities.

11 A more troublesome issue is text imposed by the computer system itself, such as the words visible on a menu bar. Is such text like the invisible stage hands of the Japanese theate—there but you don't see it? And what about text visible from another window? Should this be treated the way John Cage treated ambient sound?

12 See Larsen's *Samplers*.

it was popped up. Another approach is to consider a link name as a "prosody channel" connecting the text at either end of the link. It is typical in hypertext to assume that the reader will choose a link based on semantic or logical criteria, but in poetry there is no reason to assume prosody is any less valid as a means of choosing a link. To use the terminology we've been using throughout: bonding strength can operate through the link; bonding strength may even be the basis for choosing a link in the first place. It makes sense to speak of a "two-dimensional" prosody in assessing the relationship of prosody within the lexia to prosody through the link. Indeed, if the lexia is spatial, one may speak of a three-dimensional prosody. One point worth noting here is that the concept of bonding strength—the attraction of two text elements across a real or imagined boundary—sounds quite symmetrical, whereas most hypertext links are one-directional.[13] But the directionality of the hypertext link is not really different from the directionality of time in conventional prosody. It may be true that considering bonding strength through the link reverses the direction of attention compared to the direction of the link, but we do the same for the direction of time in assessing linear prosody.

At its most conservative, a hypertext treats the lexia as a full-fledged document in its own right; the interactive devices, such as links, may be seen merely as devices for visiting traditional documents. A more radical approach treats the episode as a virtual document. In this approach the text's center of gravity, as it were, is no longer within the lexia, but in what emerges through the use of interactive devices. At its most extreme, meaning—and syntax—are more properly a function of the episode than the lexia (Rosenberg, "Structure"). What are the implications for prosody if the episode is treated as a virtual document? This is related to a second question: what is the structure of the episode? One answer to this second question is that the episode is structured linearly by time. If we accept this idea, then prosody within the episode seems little different from prosody within the lexia, except that the reader has chosen the interactions. In the disjunctive case the reader has chosen which route to follow in operating a given acteme, and in the conjunctive case the reader has chosen the order of visiting various elements. In both cases, the reader controls how much time she spends in

13 Ted Nelson, who coined the term "hypertext," has consistently advocated that all links should be bidirectional.

any given place in the hypertext. The sense that many alternatives are possible at a given hinge point in the prosody may create the sense of that spot as a slot into which different continuations can be plugged; this very multiplicity may create a sense that some combination of some or all of the continuations is what in fact actually connects to the hinge point, which would subvert the concept of disjunctive hypertext.

But is the episode necessarily linear? I have argued elsewhere that the structure of the episode is what we make of it given the gathering interface that is available (Rosenberg, "Structure"). Alas, in most commercially available hypertext software, there is either no gathering interface at all, or it is at best extremely primitive. A gathering interface is in effect a hypertext the reader constructs of gatherings from the hypertext being read. This interface may use spatial or conjunctive methods, even if the hypertext being read uses a pure node-link model. (For an example of a commercial gathering interface operating on the World Wide Web, see Bernstein.)

HOW DOES TIME RUN IN A NON-LINEAR POEM?

Much of this paper has been concerned with a spatial approach to prosody. Yet one can hardly leave time out of the picture. The study of hypertextual time is still in its infancy. Lusebrink[14] has produced a taxonomy of time types based on narration; Calvi and Walker[15] present a hypertextual treatment of analepsis and prolepsis. These discussions, while useful, don't provide much insight for prosody. It is important to note at the outset that there are multiple concepts of time operating at once. At the most obvious level is what may be described as "usage-time," a temporality that functions like an unedited recording of what the reader actually does. In fact, such a concept of time can be misleading even in the case of very linear text. Many authors have studied "isochrony," the tendency of stressed syllables to form a regular musical beat. Even when stressed syllables do not fall according to a regular beat, the stresses themselves may so heavily

14 See Lusebrink's "The Moment in Hypertext: A Brief Lexicon of Time."

15 See Calvi's "'Lector in Rebus': The Role of the Reader and the Characteristics of Hyperreading" and Walker's "Piecing Together and Falling Apart: Finding the Story in *Afternoon*."

influence our perception of time that our sense of time may be said to be based on linguistic features like stress rather than on the purely acoustical features that would be captured by a tape recorder. Thus the stresses become our measure of time, even when their acoustical correlates do not seem to be evenly spaced.[16] Do interactive devices become the measure of time in an interactive poem? As hypertext is extended further into the fine structure of language, this may happen. Does usage-time include all the unintentional paths taken, as when one accidentally releases the mouse, or over-shoots a scroll bar?

A second concept of time is "gather-time": the time one spends constructing and reading the results in a gathering interface. As I have mentioned, most often the only gathering interface at hand is the reader's memory. Gather-time may start and stop: when a reader is foraging for an episode one may speak of gather-time as having stopped. This is no different really from the concept that the syllable-time of the poem is not running during the time it takes to find one's place in the poem on the page when momentarily interrupted. In a spatial gathering interface, is gather-time running while one changes the spatial relationship of gathered elements? Some type of time is running of course. As one manipulates gathered phrases on a screen one exists in a relationship to them that has temporal dimension. But how does that relationship map onto syllables? Is the time spent moving a phrase mapped onto all the syllables at once? Can usage-time work in this same way, given the right interface? Clearly it is possible to arrange words using graphical methods so that the eye associates all of the words together as a single object all at once, even though there may be an underlying linear structure. How does time work for such an object? There is an initial exposure time, which is arguably linear, but what about time spent contemplating the word object as a whole? What kind of time is that? Is it suspended time? Is it autonomous time, in

16 On a similar note, permit a personal anecdote. In the early seventies I made several pieces on magnetic tape using simultaneous overlays of my own voice. For one of these pieces I realized I could control these overlays very precisely by building up each fragment through making a tape loop of what was already laid down, making a tape loop of the voice to be added, then by controlling the offset of these tape loops I could get the desired effect. In one case the composition scheme called for a simultaneous attack (to use the electronic music term) of all of the voices. On one pass round the loop I felt I had nailed it exactly. But for some reason I decided to analyze the result at slow speed. Doing this it became clear that the attacks—in acoustical terms—were not simultaneous at all. What did line up simultaneously were the stressed syllables in each voice. I heard the attacks as being simultaneous—retroactive from the vantage point of having heard the stressed syllables. Linguistically the words sounded like they all started at the same time, even though acoustically this was not the case.

which the word object becomes in effect an object with its own temporality, not necessarily reconcilable with the concept of time of other objects present, much in the way two people present in the same event may not be able precisely to reconcile their individual concepts of time? Perhaps time can seem to proceed like a kind of loop, where words, having been initially examined, are treated as though they keep on playing.

Conjunctive structures bring their own set of questions to the issue of how time works. A conjunctive structure consists of all of its components resolved into a single whole. What is the time relationship among these components? It makes sense—at least metaphorically—to think of the usage-time for each component as being "equivalenced" with that of the other components. In the structures I call simultaneities, groups of words are placed in the same space, physically and logically—on top of one another. Usage history will clearly reveal the order in which the elements in the simultaneity were encountered (an order which is under some control by the user). These are different units of time; they aren't literally simultaneous, in the sense of simultaneous voices, but the term "simultaneity" is meant to convey the idea that these units of time are meant to be treated as equivalent. This concept of equivalenced time as experienced by a single user is admittedly an abstraction. Equivalenced time is a correlate of the concept of autonomous word objects—words endowed with behavior—which are so eminently possible with the use of software.

At the opposite extreme from equivalenced time are units that are completely disengaged in time, units whose time relationship to one another is completely null. Juxtaposition—bringing together elements with no structural relation between them—may be thought of as the null structure, or "structural zero," and may be considered as the most elemental maneuver at the heart of abstraction (Rosenberg, "Openings"). Clearly juxtaposition has been an important element in all of the arts for many decades. What is the null structure in the dimension(s) of time? In a hypertext, separate episodes may be time-disengaged even though the usage-time for one episode may have a clear relationship to the usage-time of another. Consider two memories, each of an incident whose time and date one cannot place, and in fact whose relative time and date one cannot place. Does it

really matter in which order the memories were recalled? The true time relationship of the memories is that they are unresolved with respect to time.

In a hypertext, time itself may become spatialized. This may occur in any number of ways. In a multimedia piece, an interactive device may permit playing a sound or movie. Such an object will have its own timeline; it is common for interactive time-based media devices to represent this timeline on the screen as a control, that the user can directly manipulate. But there is not likely to be such a timeline for the hypertext as a whole; rather the timeline for the particular media object is—in its entirety—anchored at a particular location in the hypertext. One may speak of the entire timeline as being spatialized at a particular location. Even for text, where there is no formal player object, the entirety of the text object may be anchored at a specific location. There is an important point here: for linear text, travel through the text is accomplished by reading in a linear fashion—though to be sure there are many other ways of navigating in a printed text and most acts of reading involve a mixture of linear travel along the word stream, and directly accessing various parts of the text, whether through bookmarks, tables of contents, indices, footnotes, or the like. In a hypertext, even given a linear lexia, this linearity is not likely to be used for travel. Instead, the specific interactive devices are likely to be used for travel, leaving the lexia as an anchored spot which "doesn't go anywhere." Thus to the extent there is a linear lexia, it is an anchored linearity.

MULTIUSER TIME

Throughout this whole discussion I have taken a perspective that would be called in computer jargon "single-user." We tend to view "a reading" as a single reader reading a work which has a single (even if collective) author. In the computer world, multiuser games are quite common and I feel certain that we will see an increasing number of multiuser literary works in the future. Multiuser time involves stretches of time that are not necessarily resolvable from one user to another. The events of prosody are typically passages over particular points in a poem—syllables or line breaks, etc. Where there are multiple readers in the same textual space at the same time, it may not be possible to construct any form of synchronization that would resolve the various users' interactions with the text

over time. In this sense, the concept of disengaged time is not metaphorical, but a literal description of what takes place.[17]

The questions that hypertext raises for prosody have only begun to be asked. As I've tried to show, much of our understanding of prosody has concerned the way sound events cluster when encountered in a linear sequence, and thus prosody will have to be re-thought in the context of hypertext. The central questions will include: how are we to understand prosody when clustering occurs in space instead of time? How do sound events relate across disengaged units of time? What happens to these time disengagements when the poet recites—and how indeed is a poet to perform a hypertext work?

17 It is known that the brain is a massively parallel system. A simple act of seeing involves substantial processing by each retina, even before the signals reach the brain. Is it possible that even for a single reader, the "single-user" model may not be correct? Is the brain itself perhaps "multiuser"? This is the question posed by Daniel Dennett who devised a theory of consciousness based on the concept of a parallel "gang of demons." In technical computer usage, a "daemon" is an asynchronous process—typically invisible to the user--that performs a particular type of work periodically or on request in the background. In most multiuser systems there is typically a daemon for delivering electronic mail. Another type of daemon responds to requests to view World Wide Web pages, and so on. Dennett suggests that there are centers in the brain that act as "time disengaged actors" even for a single mind. Whether or not this model of brain function prevails, hypertext is already beginning to render tangible this

Screen Dump from Diagrams Series 5 #2

VISUALITIES: THE SHAPE OF OPENINGS SPACE, BUILDINGS INVITATION, CONNECTION GUIDE SPACE

The substrate. The plane—or more elaborately the virtual plane—the arena: where the words come together. Not Olson's grid, but a continuum. (Or, on the screen, still a grid perhaps, but so fine-grained the dots nearly elude the eye: it looks like a continuum.) The stage on which words may be mounted on top of one another. (The word processor does not allow this.) And why not on top of one another: the painters put pigments atop one another, the composers play sounds at the same time, why shouldn't poets put words in the same space as well: juxtaposition, structural zero. A visuality where the juxtaposition becomes an interference pattern: grapheme jumble of full to partial unintelligibility: the lexical is annulled by multiplicity. Word cluster as a visual artifact, only partially or dimly lexical. But then: invitation: a marked spot of entry to the layers where lexicality is recovered, the words may be read as usual, the juxtaposition created lexically as well as visually. Annulment of the lexical is only the outer appearance before the cluster's opening, before the act, the dive, the tactile hand motion within the faceted word forest to that individual word tree no longer obscured by the mass of other trees—the word-part mass is a construction, what appeared from a distance as annulment of the lexical is a built simultaneity, a conjunction, phrase on phrase, a differences-music resonance chord.

The relationship between a purely visual poetry and a not at all visual poetry: where you haven't been yet. Openings not yet taken. The cluster not yet entered. Or: the construction made: resonance invoked—the after-lexicality, like an after-image, the multiplicity of phrases in the same place. The having-achieved simultaneous thought.

The supplemental channel. Where the lexicalities of different planes join, an explicit marker of structure: a syntax of making the connection direct, part to the whole if need be, folded back on itself if need be, the direct joining. Visual but a guide channel, lexicality to lexicality: a joining allowing juxtapositions equally with plain lexicalities: syntax with all slots open. The line of sight to relation mesh: visual to adjoin any achievable resonance dance, a connection guide for construction space, a pointing to the verb, a visual hand. The arrowscape skeleton latticework: arrow with all ends pointing inward: relation guide shape, thought structure trellis, skeleton of the combination act. Word combination bones, elastic, the map-stuff. The loop.

Screen Dump from Diagrams Series 5 #2

POETICS AND HYPERTEXT: WHERE ARE THE HYPERTEXT POETS?

[The following essay is a slightly edited version of an article posted to the USENET newsgroup alt.hypertext as message <506@amanue.UUCP> on 1 Dec 91, in response to a query by Shawn Fitzgerald asking if there are any hypertext writers "out there".]

I am a poet who has been working in various non-linear poetic forms since 1966. My recent work is a series in progress called Intergrams, realized as HyperCard stacks. It is not simply hypertext but uses the concept of links along with other techniques as a continuation of work I've been doing for a long time. My early non-linear work might be described as polylinear; these pieces were called word nets and consisted of interlacing skeins of writing that wandered over a two-dimensional surface. Later I became interested in how to achieve word clusters, a direct analogue of the concept in music called tone clusters that I heard my friends in the music crowd talking about. Immediately this presented a problem: how to integrate word clusters into a larger syntactical structure. By juxtaposing words together with no structural relation at all, the contribution of the word itself (as opposed to context) to defining its syntactical place was completely disrupted. This led naturally to the idea of creating a separate channel to carry syntax. Thus began in 1968 a long series of works exploring a diagrammatic syntax. As time went on I became more and more interested in the diagram concept and less interested in word clusters. Several years ago I realized that with bitmapped graphics and a mouse, I could actually implement word clusters easily: with bitmapped graphics words could be physically overlaid— juxtaposed in the same physical and logical space—and as the mouse is moved to various hot spots one phrase at a time could be brought forward so as to become readable. Intergrams combines this idea with the diagrammatic syntax. The hypertext link concept is used to navigate layers in the syntax. The concept of a diagrammatic syntax may seem unrelated to hypertext, but I will have more to say about that below.

Intergrams is forthcoming in the fall of '92 from Eastgate Systems.

Let me move on now to a broader discussion of why I think poets in particular have been slow to make use of hypertext. (I don't feel able to speak about fiction from personal experience.)

RESPONSIVENESS TO TECHNOLOGY

Some may consider this a prejudice, but I consider it self evident that the different arts respond at differing rates to technological change. The first to grab hold of new technologies always seem to be the composers. This is not surprising; all music—with the exception of a cappella singing—is inherently technological; a musical instrument is a technological device. Composers have been quick to take advantage of high technology for centuries. As Andrew Culver pointed out to me, the organ was very high technology indeed when it was new. More recently, Cage enunciated the acoustical basis of music back in the 30's before electronic music was physically possible (that has always amazed me!); at the moment electronic instruments became possible, composers were using them even though they were far beyond what an individual person could afford. (My recollection is that the Columbia-Princeton electronic music studio cost something like $80,000 back in the early fifties.) By contrast, visual artists tend to wait until a technology is poised just on the verge of being a consumer technology before jumping in. E.g. for years the only video artist in the world was Nam Jun Paik, out there all by himself; only when video was just barely too expensive to be a consumer technology did we have an explosion of video art.

Poets, alas, are the worst. We tend to wait until technology has been around for a while before jumping on it. Although poets tend not to be very conservative politically, we do tend to be conservative æsthetically, turning a blind eye to techniques that have been in use in the other arts for decades when we might make perfectly good use of them. Look at what happened in the early part of the 20th century. At almost exactly the same time Cubism turned the art world upside down and Serialism turned music upside down. The comparable movement for poetry would have been asyntactic poetry, but even today poets who have gone beyond syntax are in a tiny minority. No one considers it strange if a visual artist practices abstraction, but this is still controversial in poetry. There was of course Gertrude Stein. Robert Duncan acknowledged her influence, but still Gertrude Stein's influence on writing was nothing like that of the Cubists or

the 12-tone composers. Gertrude Stein is like a stream that goes underground. (It comes up as Cage's music! I once asked Cage if he had ever written anything about Gertrude Stein, and he answered no, that she was "too close." Pretty revealing, considering that he has written about Duchamp.) No, the writer from that period who emerges as the influential figure is Joyce, and again the impact of Joyce on writing is not comparable to the Cubists or Serialists.

But give us time, things are changing. I've been amazed that most of the poets I know now have some kind of microcomputer or other to use for word processing, even though the stereotype of poets as having meager incomes is largely true. (However! The problem of poets not being able to afford expensive machines is still having an impact, since many poets who have computers may not have high-resolution graphics.) It may take a decade or two for serious amounts of hypertext poetry to emerge, but be patient, it will come.

POETICS

Just because technology and commerce have set the table with toys does not mean that artists of whatever stripe will flock to play with them; to understand why any technology is or is not used in a particular art you must understand what is going on in that art. The post World War II era has seen a ferment in American poetics; but until recently that ferment has led in various directions that would not be conducive to hypertext. Let us look at some of the "schools" that emerged from the 50's and 60's. The three most influential groups of poets to emerge were (not in any particular order) 1) the Black Mountain poets; 2) the Beat poets; 3) the New York School. (A great number of other poetries have also flourished, but were less well organized into movements.) All three of these have had aspects that, if anything, lead away from hypertext. Projective Verse is hardly something that can be explained in a couple of sentences, but for me two of the key ideas Olson was talking about were Composition by Field (this is actually Duncan's term[1]) and a kind of real-time image concept. Composition by Field means roughly that if a continuation occurs to the poet that works locally ("within the field") it should be applied without reference to any global constraint. I suppose this could be interpreted as consistent with hypertext links, but my feeling is that in practice this tends to reinforce linearity, not encourage departure from it. The real-time-image concept is a bit more obscure; the idea is

that the time sequence of the poem should be a kind of image of the real-time act of composing the poem. Or in Olson's case, the real-time sequence of breaths should be the same time sequence as the natural prosody of the poem. Allen Ginsberg went even further in articulating a real-time concept of poetics, speaking about composing directly into a tape recorder. (He wrote a fascinating book called Improvised Poetics—probably out of print—which was the most interesting statement of poetics to appear following Olson's famous "Projective Verse" essay.) On the other hand, the "personism" of Frank O'Hara tended to deemphasize technical means altogether.

An immense cautionary tale, which looms large in the minds of many poets when alternate media, poetic forms, etc. are discussed, is the unhappy experience the poetry community at large had with Concrete Poetry. The concrete poets became prominent at roughly the same time as the movements I discussed above. Concrete poetry is difficult to describe in a sentence, but its practitioners use language elements visually, typically with the "language content"—such as syntax and semantics—completely stripped away. Many of the concrete poets were extremely polemical in pronouncing "straight" poetry (their term!) obsolete. Needless to say, this did not endear them to the rest of the poetry community. It was a tremendous tragedy. Instead of an amazing dialogue, that would have happened at roughly the same time as such ferments as Cage's Music and a whole succession of movements in art, what resulted was a total fracturing in which the concrete poets ended up talking only to themselves. To cite a specific example: there is a point in "Projective Verse" where Olson discusses the typewriter as providing a precise grid for measuring the poem. The concrete poets could have stepped into this and challenged that there is a continuum available, why be limited by a grid? A lively discussion could have resulted. It didn't. Please think about this whenever you are tempted to make bald pronouncements to writers that hypertext is the future of language. That may indeed be the case, but it is a difficult challenge to educate writers to the wealth of new techniques that technology affords without disparaging the work they are now doing.

In recent years the emergence of the so-called Language Poets has changed the situation in poetics considerably. (My feeling that the term "Language Poets" must be prefaced by a distancing qualification is painful; I know many of these

people, and it is arguably "my scene". But I cannot in good conscience condone the hijacking of the word `language' by a small group of poets, nor can I fathom the seeming acquiescence in this of the poetry world at large. The term has long since become an accomplished fact; still one must protest!) This movement has grown from the work of early pioneers Clark Coolidge and Jackson Mac Low. I will not propose to give a definition of it, but there are many aspects in the poetics of these poets that make it much more likely for hypertext to be accepted.

I do have a prejudice that the best art will come from a situation where the artist is preadapted to such concepts as hypertext by an ongoing æsthetic. It is not so simple for someone used to writing one-word-followed-by-the-next designed to be read the same way to be suddenly handed hypertext and to make good use of it. On the other hand, writers who have been using various other kinds of experimental techniques that may seem unrelated to hypertext are much better adapted to the non-linearity of hypertext. If you would seek out hypertext writing, seek out experimental writers in general; if they are not currently working with hypertext, offer it to them.

I would like to now go into some detail about specific elements of poetics that I believe would predispose a poet to being receptive to hypertext.

PRECOMPOSITION

By precomposition I mean composing a work in layers, where work in one layer may affect the entirety of the succeeding layer. Precomposition is widely practiced in music; this is one area where visual artists have taken the lead and have been doing precomposition for centuries. At one time making a painting meant scouring the woods for materials to grind into pigments, then building a support, applying the priming, applying various undercoats—all before a single square millimeter of the final surface of paint was applied. Despite the great heterogeneity of American poetry, there has been almost a conspiracy of romanticism about the act of composing itself that has led poets to look unfavorably on precomposition. Many of the language poets do practice precomposition, and consider language as material in much the same way that a visual artist might consider paint as a material. An important part of the hypertext act is the linking of pre-existing materials. This does not sit well with a poetics that emphasizes real-time poetics or the spontaneity of the act of composition; on the other hand it agrees very well indeed with an established practice of precomposition.

INDETERMINACY

A hypertext is a fundamentally indeterminate work: the user is totally free to choose whether or not to follow a link, and the node sequence experience of each "reader" (we really do need a new verb here) will be different. Cage has opened music very wide indeed to indeterminacy; many dozens of composers have used various forms of aleatory methods, even though they may not agree with other parts of Cage's æsthetic. One might even argue that the improvisational bent of classical jazz is a preadapting form of indeterminacy. While Jackson Mac Low stands out as having used various forms of indeterminacy for decades (he is a long-time friend and associate of Cage, and began using chance at about the same time as Cage did), the use of indeterminacy in poetry is not common. Indeterminacy in music tends to happen in the degree of control or choice in the instructions given to other people. (The musicians.) Poets are not used to giving instructions. Composition is a solo act, performance is a solo act. (Even in group readings, which are serially solo.) There have existed such traditions as the Japanese Renga, in which several poets would each in turn produce a line, but poems for multiple voices are not common. This is one area where there may be no obvious preadapting poetics; exposure of poets to indeterminacy is as likely to come from confrontation with hypertext itself as from anything I know of that is happening within poetry.

NON-LINEAR STRUCTURE

Perhaps this one is so obvious that it hardly needs mentioning. This is an area where the devastation caused by concrete poetry is particularly painful. There is no natural temptation for poets to experiment with diagrams or other forms of non-linear structure; even many of the language poets are amazingly conservative in the cosmetic "look and feel" of the poem. The difficulty of giving up linearity is something about which I can speak personally, having been doing non-linear work for 25 years. There is no doubt about it: doing non-linearly structured writing is a leap, an immense leap. That writers are not jumping in en masse is hardly surprising. I fear it will take an entire generation before it is as common as the advocates of hypertext would like.

THE AWFUL ASYMMETRY OF HYPERTEXT

"Traditional" hypertext entails imposing a non-linear structure on a locally linear substrate. Face it: this is an inherent asymmetry. The act of continuing to read on is simply nothing like the act of following a link. This may be of little concern to medical students making use of a hypertext medical information base, but to writers, who have an æsthetic, it is a major concern indeed. Batting back and forth between ordinary linear reading and following links is a kind of bending back and forth that can produce the æsthetic equivalent of metal fatigue.

There is one obvious cure for this problem: the concept of hypertext must be extended all the way into the fine structure of language itself, at least as far as syntax. This is an extreme position, a radical position. This is the place where my diagram poems began, more than 20 years ago; the atmosphere up there is so rare, the concept is so radical, that I have had to climb down from there to produce works that other people can understand. Still: if you truly believe in hypertext then you must be prepared for those who believe that we have not arrived at hypertext until we begin forming individual thoughts as hypertexts, thoughts that cannot be expressed in any other way. The hypertext link-following mechanism can be used to carry the actual syntactical relationships of thought. A diagram syntax is the first step in this process. There is no fundamental difference between diagrams and hypertext, except that a hypertext is a diagram where the diagrammatic links go off the edge of the page. Or to put it slightly differently, a hypertext is a virtual diagram with an unbounded virtual page.

One is somewhat daunted in making such proposals by the mass of evidence concerning the degree to which some parts of our language capability may be biologically based. Like a lot of poets, I am sure, I initially found the idea that the brain is biologically wired for syntax to be distasteful. Bearing in mind that any mixing of biology and the social behavior of human beings is potential social dynamite, still, the amount of evidence for some degree of biological basis for language is now nearly overwhelming. If the brain is indeed wired for syntax as we know it, how are we to achieve this fine-grained hypertext, this hypertext syntax? It will be difficult. Still, no one would deny that the conformation of our bone structure is largely biological, but that has not abolished the dance. To the contrary: we appreciate all the more keenly the work of a dancer who expands our ideas of what we thought the human body could do. If syntax is biological, it

makes all the more valuable artistic work that will take us beyond this. That it is so difficult is part of what makes it so alluring.

To be candid, I expect few takers along this road, at least for a while. There is another method of dealing with the awful asymmetry of hypertext which I suspect that poets in particular will find much more congenial. We are accustomed to taking our word journeys on the sentence rail. By composing word groups that are static, that "don't journey", the way is open to using other forms of travel to navigate from one word group to another. The proposal above entails abandoning traditional syntax and using hypertext itself as the syntactical channel; this proposal would keep syntax locally but abandon the sentence as the vehicle for word travel. All travel would be by hypertext link. A poet would obviously be preadapted for this idea by a more centered kind of writing, possibly even a minimalist kind. Again, we are talking here about a specific æsthetic bias, one which it is unreasonable to expect of any but a minority of poets, at least at first. (My current work uses both the more radical kind of syntax and this second idea.)

POETIC CULTURE

The culture of contemporary American poetry is an oral one. Those who care about their poetry expect to be able to hear the poet recite it. Believe me: to have work which cannot be recited—because of its intrinsic structure—is a severe handicap. It means virtually renouncing "the scene." Asking a poet to do this is asking a lot. I believe very fervently that McLuhan had it exactly backwards, that speech is inherently linear, by virtue of the real-time synchronization constraint between speaker and listener, whereas writing, by allowing for diagrams, has the inherent capacity to be non-linear. Perhaps the single most exasperating thing I have faced in my career as a poet is the reaction to inherently non-linear work: "Well, I can't believe there's not some way to perform this!" I get this all the time, from people who should know better. The urge is just too strong, the sense that the community comes together at readings is too overpowering for most poets to face the idea of work that is unrecitable. I believe that this is one area where only the technology itself will improve the situation. Yes, this is an area where the lack of venue for non-linear writing hurts. It will take the emergence of a network community of writers to overcome this problem. And when it does come about, there will be an acute problem of making sure there aren't two communities, separated from one another, the networkers ghettoized from the oral poets. Again,

it is CRUCIAL that those of us who are experimenting with new technologies in writing do our utmost to avoid anything that even smacks of hinting that there is something somehow "wrong" or "archaic" about those pursuing traditional media. We experimentalists need the traditionalists—without the tradition, how can there be an experiment? But the converse is not true. The traditionalists can go on quite happily without us. The risk that hypertext writers will become ghettoized is very great. It will take great skill to keep the lines of communication open.

[1]It seems this was an error on my part, and that the term 'Composition by Field' is indeed due to Olson. [jr, 11/19/2012]

NOTES TOWARD A NON-LINEAR PROSODY OF SPACE

LINEAR METRICS

BACKGROUND: DIMENSIONS OF INEQUALITY AMONG SYLLABLES

The fundamental axiom of prosody in English is that all syllables are not created equal. (This is not the case in all languages; in Japanese, apparently, syllables are considered equal enough, metrically, that they are simply counted rather than being distinguished by structural patterns.) By plotting the time sequence of how these inequalities occur one obtains metric—a double plural, meaning a multitude of metrics. The most well known is the stress degree metric, which forms most of the discussion below, but there are several others. (1) Pitch Degrees. Although I have not analyzed this in detail, there is a prosody at work in many American communities, most notably in a style of reading in the black community, in which I believe pitch degrees play a more important role than stress degrees. This forms the basis of an incredibly rich prosody, a different prosody than the classical one, which must be analyzed on its own terms. Whereas stress is a linguistic property, pitch is a purely acoustical property; thus a pitch degree prosody is more absolutely musical—in a literal sense—than a stress degree prosody. A pitch degree prosody is more free to use an absolute musical sense of time and structure; the sound structure balance is different—not all phonemes carry pitch. Where a pitch degree prosody and a stress degree prosody are present at the same time, incredibly subtle effects are possible. Pitch degree prosody deserves extensive analysis in its own right.

(2) Vowel Position Degrees. Again, my impressions here are based on intuition rather than detailed analysis. My understanding of this goes back to a somewhat delphic comment by Robert Duncan. Duncan spoke often of "tone leading vowels" in talking about Pound. Not understanding what he meant by this, but wanting to know more, at a reading once I asked him where in Pound's writings

I could find the discussion of tone leading vowels. To my total astonishment, he replied that it wasn't anywhere; Pound had used the term in letters to Duncan, simply assuming that Duncan would understand what he meant; Duncan was left to figure it out for himself. Hearing this I simply couldn't resist: "Well, what is it??" Duncan said two things, one of which was quite straightforward, the other of which was extremely obscure. He explained that when a diphthong (a glide from one "pure vowel" to another) occurs, the "leading tone", i.e. the pure vowel at which the glide begins, plays a special role in terms of later reinforcement. So far so good. Then he said: "When you hear a sound, it's reinforced when you hear it again. But it can also be reinforced when you don't hear it again." Just as Duncan was left to figure out on his own what Pound meant by "tone leading vowels", I felt strongly that I should simply accept this remark as a gift and work on my own to figure out what it meant; I didn't ask anything else.

I certainly don't wish to be in the business of interpreting Robert Duncan, even less interpreting Duncan interpreting Pound; but the train of thought that started with Duncan's remark has ended up with the idea of a vowel position degree metric. Vowels are sometimes diagrammed as a trapezoid which plots the position of the tongue in the mouth. I believe (again this is based on intuition, not analysis) that the possible vowel positions tend to cluster into just a few groupings, and that the plot of which grouping occurs at a given moment forms a metric. This metric is probably superimposed on the stress degree metric, again giving a very fluid effect. (I don't know if this has anything at all to do with "tone leading vowels", but I certainly do hear something like this frequently when I read or listen to Pound.)

THE CLASSICAL STRESS-DEGREE METRIC STEREOTYPE

Before giving my own approach to the stress-degree metric, let me briefly sketch the stereotype most people have for how this works. If you want to argue that the stereotype is not accurate in scanning this or that poet, I would hardly disagree; but I think this stereotype does have a lot to do with attitudes towards metrics. Hopefully, by presenting a different approach, the stereotype can be broken.

There exist -- a priori, in advance of any particular poems -- a collection of abstract patterns of stress degrees. These patterns may be called templates. They tend to have names, e.g. iamb (‘|), trochee (|’), anapest (|”) etc. (I’m using ‘ for an unstressed syllable, | for a stress.) Scanning is an activity consisting of pattern-matching template instances against the words of the poem. In the most formalized case, we say a line “scans” if repeated instances of the same template, end to end, match against the words of the line. There is a tendency to speak of poets “writing in” a meter, which means choosing a template in advance and writing lines which will scan to that template.

Readers familiar with prosody terminology may note with some surprise that so far I have not mentioned the word ‘foot’. Unfortunately, the word ‘foot’ is deeply ambiguous: It has at least two wildly different uses in prosody. Because I believe this plays a crucial role in the widespread prejudice against and misunderstanding of metrics by many people, I want to explain this in detail, and will propose a different terminology to avoid the problem. Meanings of the word ‘foot’: (1) ‘Foot’ is often used to designate a type of template. (E.g. one speaks of “an iamb”.) As such it is part of an a priori scheme, given in advance of the poem. (2) ‘Foot’ is used to designate the group of syllables in the actual poem which is matched by one of the templates. In this usage, a foot is a unit of meter: just above the syllable, which is the atomic unit for metrics, and below the metric line. (Of course this is not meant to deny that a foot can contain just one syllable or a line one foot.) I.e. here foot is a living breathing part of the poem, with a specific tangible identity in a metrical structure.

Many different strains in poetics, such as the Open Forms of the Black Mountain poets and Improvisational Composition of the beats, have led in directions away from the a priori, away from choosing a fixed form given in advance. Because the concept of ‘foot’ is so closely linked with a system of a priori templates, this has led to a sad neglect of metrics generally, and to interest in the concept of mapping the metrical structure of living breathing poems specifically. To avoid this problem, in what follows I will use the term measure for a unit of meter intermediate between the syllable and the line where we make no assumption in advance for what shape a measure may take.

BONDING STRENGTH

I wish to propose yet another dimension along which syllables can vary: bonding strength. Actually bonding strength applies not to syllables but to the boundaries between adjacent syllables. Before going into the details of how this works, it is important to note two important points about this methodology. (1) The methodology applies against the sound of the poem. In some cases this may be difficult to infer from a printed text absent a live or recorded reading by the poet. This idea will be quite threatening to those academics who believe that all there is to know about a poem is contained in the printed text. However, there is nothing terribly original in the idea that prosody means looking at the structure of the sound: it was Ezra Pound who said if you want to know about the prosody, open up your ears and listen to the sound it makes. (2) The methodology -- and I emphasize that it is a methodology, not a theory! -- is ruthlessly empirical. It works by examining a recitation and trying to answer the question: what is the metrical structure of this recitation. It may or may not yield useful results for any given recitation. It is not designed to answer what may be a fallacious question -- what meter did the poet "write in" -- but simply to discover whatever metrical structure happens to be there. Because poets build a voice, there is a reasonable chance that in many cases conclusions can be drawn about metrical structure, but perhaps not.

Each syllable has a bonding strength for the syllable before it and after it. By bonding strength is meant: (artificially) inject a pause at the syllable boundary in question, and then judge how natural the pause is against the recitation. There will be different degrees of naturalness -- different degrees of resistance to the injection of pause -- at different syllable boundaries. We say a syllable boundary has a high bonding strength if an injected pause is extremely unnatural compared to the recitation; where the pause is natural (or already there, of course) bonding strength is low. I will call a measure a string of syllables bounded on either side by low bonding strength and having only high bonding strength in any internal syllable boundaries. It is this empirical, overtly sound-based concept of measure which I wish to use as a replacement for usage (2) of 'foot' above.

Of course this will work with any recitation; it does not have to be poetry. I propose that the speech rhythm of English prose tends toward what I will call The Standard Measure, defined by the following rules: (1) A measure has only one major stress; (2) the stress tends to come at the end of the measure, but (3) if there are only unstressed syllables following the major stress out to the end of a natural grammatical boundary, those syllables will be incorporated into the measure. (You can think of this as an attempt to reformulate the classical idea that English is "predominantly iambic" while institutionalizing the counter-examples.) The concept of "grammatical boundary" is really the same concept of bonding strength applied to syntax. Although I haven't analyzed this in detail, a simple explanation would relate grammatical bonding strength to the distance that must be traveled in a parse tree to get the closest common antecedent to two consecutive words.

This is not to say that "the standard measure" will in fact always occur. Robert Creeley, for instance, is well known for having many non-standard measure breaks in his poems, and many other poets use non-standard measure breaks quite prominently. Detailed analysis can produce some interesting results. For instance, in some of Creeley's poems, lines are predominantly either one or two measures; interestingly in the two-measure lines the internal measure boundary is a standard measure break. (I.e. the celebrated Creeley line-break is exactly that, a line-break, not a measure- break.) The non-standard measure boundaries are easy to hear, but the internal standard measure boundaries do not stand out so prominently. (Of course you would "hear their absence" as a lifeless regularity.) In these poems one could say that the grammatical structure goes from the middle of one line to the middle of the next; but the line endings are articulated sonically by the prominent non-standard measure boundaries. This is a familiar metrical structure going back hundreds of years, all the way to Anglo-Saxon poetry.

NON-LINEAR METRICS

Prosody is traditionally the study of poetic sound structure as mapped in time. Because of the inherent linearity of time, this poses an obvious problem for non-linear poetry. Shall we say that prosody only applies within the lexia, thus

punting the problem completely? Some may take this view, but I find it personally distasteful. It exempts from prosody details of hypertext structure which I think clearly need to be considered. In Michael Joyce's work, the names of links are so clearly musical it takes one's breath away: they are part of the prosody. Given that I am on record as advocating taking hypertext into the fine structure of language, thereby fragmenting the lexia, I simply can't accept leaving prosody as an inherently linear concept that applies only inside of lexia.

It would be nice if what follows were as well worked out as what preceded, but at this point I have only questions and some launching points for a view of non-linear prosody.

LINEAR PROSODY AS A PROSODY OF SPACE

There are some concepts in prosody that are so overtly time-based that one simply has to give up on them in non-linear poetry. Isochrony -- the tendency of major stresses to fall in an even musical beat -- is one such concept. There may be isochrony within the lexia, but given that outside the lexia there is no concept of time, in non-linear prosody one can only treat isochrony locally. However: In a one-dimensional structure, time and space are very nearly the same thing. What may appear to be a time-based concept may in fact be a space-based concept. The concept of bonding strength, as articulated above, occurs between adjacent syllables. 'Adjacent' is clearly a spatial concept, not a time-based concept. I defined bonding strength as the resistance by a recitation to the injection of an artificial pause -- time language again -- but one could just as well speak of injecting space into the poem as time. Thus:

SPATIAL BONDING STRENGTH

The concept of bonding strength works in any kind of topology. One may speak of the bonding strength of two adjacent units as their resistance to the injection of artificial space. The replacement for the concept of measure above is a spatial clustering. This has some interesting consequences for "traditional"

hypertext rhetoric. Whereas the classical hypertext link is typically discussed using travel vocabulary, a spatial prosody would ask the question:

What is the bonding strength THROUGH the link?

(This has an interesting resonance with the Kaplan/Moulthrop concept of hypertext warping space.) It may be objected that in trying to assess bonding strength through the link, one is reversing the direction of the arrow (assuming a one-directional link, and recalling Ted Nelson's caution that all links should really be bidirectional). Well, having no problem measuring bonding strength "against" the arrow of time in linear poetry, I have no problem doing the same thing "against" the arrow of the link. Travel may be the right vocabulary to use, but one can also speak of the attraction of two nodes for one another; yes, there may be an asymmetry in link direction, but there is always an asymmetry between "located here" and "could be located there". In discussing hypertext there is an overt tendency to discuss following links based on similarity to where one is; but perhaps not! Perhaps one wants to take a link based on dissimilarity to where one is! (Robert Duncan: "A sound is reinforced when it is not heard again ...")

Prosody may form an overt basis for following links. I find myself following links in Michael Joyce's work based on sound structure all the time. Perhaps I'm "not supposed to do that", but I do, and without having asked him about this I think Michael wouldn't mind that I do.

Prosody must work THROUGH links, not just inside "the lexia" (if one has lexia ...)

I speak about "links" here only because that is the most familiar form of hypertext structure; this discussion generalizes to conjunctive structures as well.

AGENCIES

By "agency" I mean a unit of doing. Because there is an overt structure resulting from hypertext linkage, we have a bit of a tendency to focus on that structure rather than the structure of what the reader does. Consider the familiar

node/link/lexia model. A reader follows a series of links, then -- for whatever reason -- decides to backtrack. That series of nodes is a unit of doing. Unfortunately, most of the software does not treat it that way. For instance I can't save the series. Perhaps I can save my entire history, but I typically can't mark it to denote the way I think the links clump together. Perhaps the best I may be able to do is save my history to a text file and then import it into some tool where I can annotate it.

Doing has boundaries, has a shape, has units. The structure of doing unfolds against the skeleton of the hypertext structure, but it may not be the same. The hypertext structure may appear to be disjunctive, but the reader can do conjunctive things with it anyway. (E.g.: "These are the links I really like, they seem to work together ...) In a hypertext with conventional lexia, there is a concept of bonding strength among lexia.

When I speak of agency, I am not speaking of history. Agency and history are related but are not the same. The structure of history is relentlessly linear: "this is the sequence of what happened ..."; the structure of agency is likely not to be linear. There will be clusters of doing that include the same event. Some agencies will stand out clearly in the reader's mind; these may clump together even though they are isolated in time by navigations that don't stand out in the reader's mind as a clear agency: agencies don't always happen.

THE GRANULARITY OF AGENCY

We have, perhaps, come far afield of prosody. I am propounding a view of linear prosody as clustering in one-dimensional space; from this one generalizes to clustering in a space of arbitrary topology. The granularity of this clustering as it traditionally affects prosody is very fine: down to the level of the syllable. A hypertext prosody granularity of space can become this fine only when hypertext is taken into the fine structure of language. I am always (it seems) advocating for this, but with few takers at present. With a more "traditional" lexia, there is a "granularity boundary" concerning units of prosody that are inside the lexia vs. those that transcend the lexia. Is the poetic lexia a stanza? Is it a line? Is it yet larger than the stanza? Even I would balk at taking hypertext inside the word.

The words are given to us, by and large; it does not seem reasonable to me to intervene in that natural process with an external administration of hypertext structure. It should be noted that in the discussion of linear prosody above, the "standard" speech rhythm of English does not interpose measure boundaries inside the word either. If I will not intervene inside the word with hypertext structure, and the natural rhythm of English does not intervene inside the word, that leaves the entire measure in a hypertext context likely to be linear. At its finest granularity, hypertext structure relates measures.

In the large, the structure of agency extends to the session. While surely multiple sessions will associate in the mind of the reader, just as there is a natural limit at the measure boundary -- even though measures contain still smaller units, namely syllables -- there is a natural boundary at the large end of the granularity of agency at the session. You have a well defined cut in doing when you sign off, surely. As obvious as this sounds, there is still a great deal of work to be done in understanding what the identity of the session should be. Hypertexts tend often to be large; it can take many sessions before the reader even begins to get a feel for how a particular hypertext "is supposed to work". Because the number of paths is genuinely infinite, the reader may have no help at all in deciding when a session should end.

Shall we say as writers, overtly and explicitly, that we want the reader to compose a structure of agency? Shall we say that even if the writer does not want the reader to do this, readers will want to anyway?

It seems the prosody of hypertext needs new terminology. I am comfortable with the terms syllable/word/measure/lexia, but up from there the familiar terminology breaks down completely.

MULTIPLICITY OF AGENCY

There are clearly multiple concurrent agencies in a multiuser hypertext that is being used by several users at once. An interesting rhetorical question arises from this: do we need to discuss multiplicity of agency in the context of a single reader reading a single hypertext? Put differently this is the question: How

does a single mind apprehend a network? Must one simulate multiple concurrent agencies in one's own mind? (It is interesting to note in this regard that some theorists of mind -- notably Dennett -- have proposed that despite the illusion that the mind is a singular stream of agency, in fact the mind is a multiplicity of concurrent agencies.)

Many in the hypertext community will surely object that in this whole discussion so far, I have spoken as if reader and writer are different people -- i.e. the hypertext is "closed". It seems reasonable to me to entertain the idea that writing and reading are separate agencies even if done by the same person in a single session. Writing-as-reading simply becomes yet another dimension to the structure of agency.

Concluding Question:

How do we -- or do we at all -- code the hypertext for the structure of agency?

THE WORD THE PLAY ATTACHING AT A WIDE INTERVAL

"Water bird" —Washoe

Combination from a fixed palpable set given in advance: the word is composite, an object with parts from a small number of choices: syllable, phoneme, letter of the alphabet, the word is an assembly of sounds, of letters, a play of referenceless tokens into breathable sign object (q.v. Derrida): the familiar narrow interval. In composition by vocabulary the interval is wider: a fixed set of words, composed in advance of the poem where the word is atomic, a moveable catalysis object (sign stuff dissolved, a resonance token indivisible but with valence): the composite is wider than the word, words are brought together as a kind of super-word, phrase as self-mobile word molecule, divisible and yet acting only as the clump. The units shift up one notch in scale, phrase acting in syntax as the word does normally. Vocabulary, word-set: the pieces set out for play as beings are, whole, alive, not to be dismembered. A play at a wider interval: the word not as supplement-of but as activator of resonance: like the enzyme, it disappears having merely activated an energy reaction of presences one possesses already. The word is not a gift, is not an epiphany, only a bending of the resonance space nearby. The real stuff of poetry is still human lives, as it always has been, notwithstanding that the play occurs, screenwise, in a machine.

That the super-word holds, can chord-achieve its precipitation of the oracle key: the word is reencountered as a member of the word set with a chance present location: the word is new, new to this place, newer than the non-set "spontaneously voluntary" word we hardly touch: the external word, the word that originates externally. The word is always there before we get there, and yet for the poet, traditionally the word has always seemed to originate "inside". The links of words into super-word, that is what is now internal. But having been made, in the machine, the super-word can be made external as well: the flow in and out, like sodium channels that power the nerve force. The reaching for the word becomes nearly physical, as if pulling it from its place on a signboard. Make the word new by putting it outside, bringing it inside again as super-word, then put-

ting it outside again: like breathing. How crystalline the word gets, there outside the mind, doing battle with erosion. Computer screen as crystal garden, part of the mind / not part of the mind.

Metamorphosis, migration, evolution, reassembly: substitution may occur at the play of super-words one for another, bonds broken and reattached elsewhere in the mixing of the atoms at higher energy, a fractally point-sharp force of surprise. The phrase rejected in editing becomes compost, food: components as nutrient put back for recombination in newly attached super-words, sampled and ordered by chance: the mixing, recombination—~~as~~ in life—~~t~~he source of energy, of diversity, a flow cycle reaction arena, a self-sustaining word process ...

MAKING WAY FOR MAKING WAY:

Co-striation Act Topographer of the Mingle Scriptor Transform Dance

A review of

Of Two Minds: Hypertext Pedagogy and Poetics by Michael Joyce

The University of Michigan Press, Ann Arbor, 1995, 277 pages, ISBN 0-472-09578-1

The palpable sense of handling -- as distinct physical beads -- the crown jewels: that's what it felt like to prepare for this review by transcribing (by hand, using Graffiti handwriting software on a Newton) and then manipulating exact quotes from this continuous song of a book by Michael Joyce. Though surely there are in this amazing book some points to which one might object, I find myself standing in wonder at the breadth of riches it contains. A definitive collection of essays dating back to 1987 by the man universally acknowledged as the father of hypertext fiction, this book is cause for celebration by anyone with an interest in literary hypertext. The book is organized into three parts: 1 Of Two Minds: Hypertext Contexts; 2 Siren Shapes: Hypertext Pedagogy; and 3 Contours: Hypertext Poetics; but that list hardly begins to suggest the breadth of topics covered. I found my list of saved quotations aggregating around the topics Programming, Polemic, The Sheer Joyous Images, Exploratory vs. Constructive Hypertext, Contour, Pedagogy, Spatiality, Meaning, Reality (Virtual and Otherwise), Time, Mind -- and even this list seems to me hideously incomplete. An entire essay (or several!) could easily be devoted to discussing just the topic of exploratory vs. constructive hypertext, which is an important recurring theme in most of these essays: "Scriptors use constructive hypertexts to develop a body of information that they map according to their needs, their interests, and the transformations they discover as they invent, gather, and act upon that information. More than with exploratory hypertexts, constructive hypertexts require a capability to act to create, change, and recover particular encounters within the developing body of knowledge." [emphasis mine].

I was extremely intrigued by the section on pedagogy. Michael Joyce's credentials as a postmodernist are impeccable, and yet here you will find detailed protocols for the use of hypertext in the classroom to teach what many in the

literary profession would shudder to have to do: remedial composition. If anyone believes hypertext writing is just for eggheads after reading part 2 of this book then I give up, such people are not reachable. That is not to say this section of the book is any the less literary:

"We redeem history when we put structure under question in the ways that narrative, hypertext, and teaching each do in their essence. Narrative is the series of individual questions that marginalize accepted order and thus enact history. Hypertext links are no less than the trace of such questions, a conversation with structure. So too, the networked classroom is a place of 'making do' as constructive action. All three -- narrative, hypertext, and classroom -- are authentically concerned with consciousness rather than information, with creating knowledge rather than the mere ordering or inventory of the known. The value produced by the readers of hypertexts or by our colearners is constrained by systems that refuse them the centrality of their authorship. What is at risk is both mind and history."

There are a few places in this book where I find myself taking distinct objection. In places the degree of polemic does reach the level of being an unguided missile: "The book is an obscure pleasure like the opera or cigarettes. The book is dead -- long live the book." I know many hypertext writers feel this way, but I happen to feel we are not doing our cause any good by this kind of talk. (It is worth noting that the "obscure pleasure" of this book extends to the binding: it's sewn!) More serious is the issue of programming -- or to be more accurate, the absence of a serious consideration of programming. We hear frequently that "Hypertexts are read where they are written and written as they are read." True enough -- but what about programming? To this reviewer, omitting the programmer from a discussion of writer/reader duality is to miss the essential fact that hypertext involves a triad of reader/writer/programmer; to exclude the programmer from the reader/writer rhetoric is as serious a problem as excluding the reader-as-writer from a discussion of reading. Example: "How, in the landscape of the city of text, can the reader know that what she builds will move the course of the river?" She could do worse than to start by reading the code. (If the code is not accessible then that itself is an important issue for hypertext rhetoric.) Or: "Hypertext programming -- in the double sense of instructed machine behavior and information content:

the thing that, like situation comedy, weather map, or docudrama, is shown on the screen -- is privileged, centralized, and self-sufficient." Privileged? Self-sufficient? How incongruous this sounds amidst the many pleas for constructive over exploratory hypertexts. The degree to which hypertext software systems should be extensible -- precisely so the programming is not "privileged" -- is contentious; suffice it to say that this reviewer believes that systems must be extensible.

Surely the most challenging enterprise in reading this book -- and the one most likely of highest priority for those in the hypertext research community -- is deciphering Joyce's absolutely essential concept of contour. (One notes with a wry smile that at one point Joyce himself pleads guilty to writing with obscurity on this subject.) The contour is: An intersection in pathways through the many lexia / An outline, a memory, a resonance in the replacement of one writing by another / A density from following multiple links / A contour line, an altitude, an ache of criss-cross as we stretch the end-to-end link iteration from one heart-foot left behind -- these are guesses on my part. The contour is: a multiplicity. As an extremely young poet, I, like many others of my generation faced a kind of rite of passage in poetics in wrestling Charles Olson's famous essay "Projective Verse" into some form of ready raw-material congeniality; it was something no one could do for you. So too to understand Michael Joyce -- another breather in that being-sigh of reception from Olson, the archeologist of morning -- when he speaks about the contour is to undertake, not to be handed a cute package. And yet, one is convinced -- haunted perhaps -- this is indeed the heart of the matter: the topographic after-image in the face of the constructive act formed of so many individual bits of interactive doing.

Those of us who write interactive literature believe, down to our very bones, that a literary hypertext is what it is through the interactive choices of the reader; if a few simple sentences strung end to end would serve just as well then we would put those into play instead. We put in play text interlaced with interactive structure operations because we believe nothing else will serve the purpose. Given that, it is worth asking whether we can expect that something as complex as the mechanism by which the reader arrives at a resonance shape in how a hypertext means could possibly be formulated in a few simple sentences of prose. A hypertext is a journeying system; in asking an author to do more than

set out beaconing launch points toward our understanding of how the energy in a hypertext works is perhaps asking too much. To ask the reader to make do with a neatly formulated few simple sentences of guidance is perhaps asking too little.

I close where I began, with the last and first word about Michael Joyce: that he writes so excruciatingly well, those of us who follow can only read such words as, “We are the children of the aleatory convergence. Our longing for multiplicity and simultaneity seems upon reflection an ancient one, the sole center of the whirlwind, the one silence” and murmur. We nod: oh yes, oh yes

THE INTERACTIVE DIAGRAM SENTENCE: HYPERTEXT AS A MEDIUM OF THOUGHT

DIAGRAMS: A SEPARATE CHANNEL FOR SYNTAX

To begin with the elemental, the "structural zero," juxtaposition: the act of simply putting an element on top of another, with no other structural relation between the two elements except that they are brought together, is the most basic structural act, the most fundamental micromaneuver at the heart of all abstraction. But consider the problem of the poet in bringing this about. When a sound is played simultaneously with another sound, the result is a sound. When a painter places a bit of colored space on top of another bit of colored space, the result is a bit of colored space. A mathematician would say that the domains of the composer or visual artist are closed with respect to the operation of juxtaposition: the result of juxtaposing two elements from the domain is another element from the domain. But what happens when we juxtapose words? Whether it is done by means of sound -- either via simultaneous readings by multiple performers, or by overlaying magnetic or digital media -- or visually, the result of juxtaposing words -- in the almost palpable physical sense of putting them directly on top of one another -- is likely to be sheer unintelligibility: one will be lucky to make out any of the words at all. How is the poet to achieve juxtaposition with no sacrifice of intelligibility?

But it gets worse: how can direct juxtapositions of words be used in larger structures? It is not hard to work in modes that give up such structures as syntax. One simply does without. Asyntactic poetry is a large and fruitful domain in which to work. On the other hand, giving up all possibility of structure is giving up a great deal indeed. Syntax is at the heart of how we normally structure words. How does one achieve such structuring and yet still have complete freedom to use juxtaposition wherever it is artistically important? How does one designate the structural role of a juxtaposition in a larger structure? One could put this question a bit more crudely by asking: What is the part of speech of a juxtaposition? The composer John Cage once criticized the twelve-tone system as having

no zero[18]. One could say that syntax "has no zero": in a sentence every element has its structural role with respect to the syntax diagram, or parse tree; there is no way to have words in a sentence whose syntactical relationship to one another is the null relationship: nothing at all except that they are brought together. How can the poet have her cake and eat it too? I.e. how can one keep both syntactical null relationships and much more elaborate relationships, in which juxtapositions act as elements?

These are some of the formal problems that have motivated my work going back more than 25 years. A method for approaching the second problem -- how to incorporate null structures as structural elements -- became apparent long before I realized how juxtaposition could actually be implemented. By devising an explicit visual structural vocabulary -- separating syntax out into its own channel, so to speak -- structural roles could simply be directly indicated. The elements occupying those roles might be words or word clusters or other structural complexes. Thus began a long series of works called Diagram Poems.

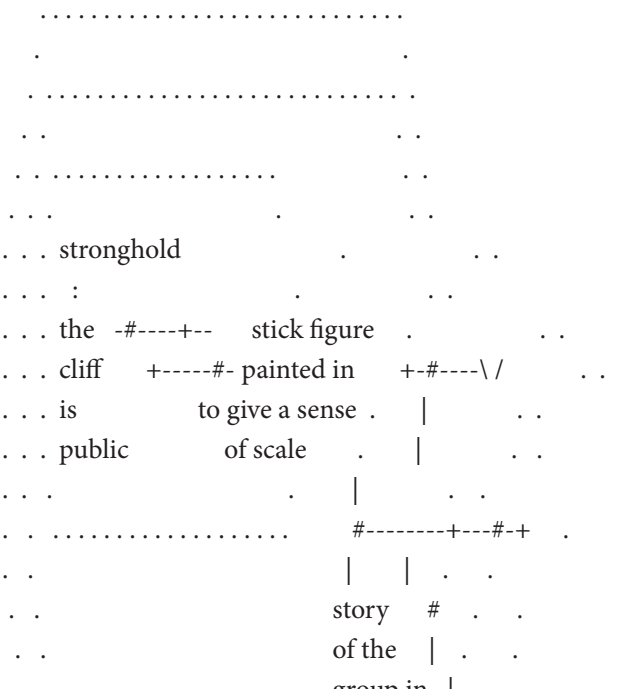

18 See Cage's *Silence*.

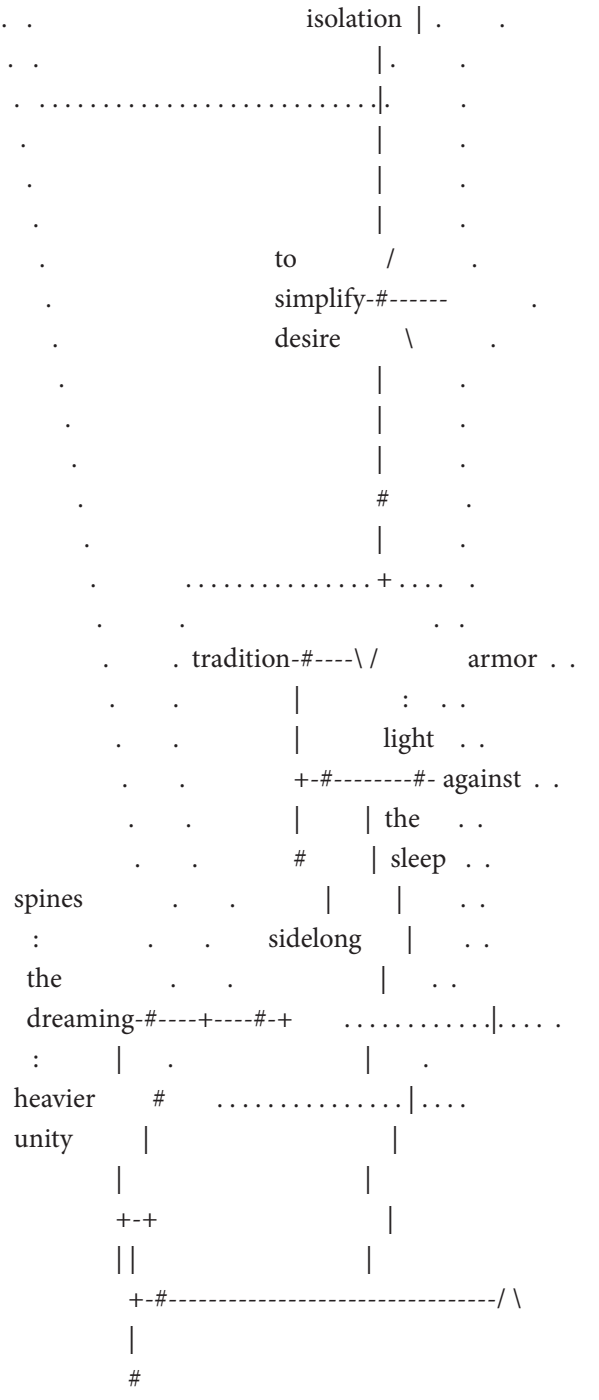
isolation
to
simplify-#------
desire
tradition-#----\ /
armor
light
against
the
sleep
spines
sidelong
dreaming-#----+----#-+
heavier
unity
+-#----------/ \

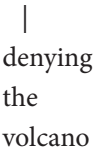

Figure 1: A diagram poem from Diagrams Series 3

Figure 1 shows a poem from Diagrams Series 32. It illustrates many of the facilities provided by the diagram notation in a variety of works spanning a large number of years. The configuration:

```
x -#----+--
     +-----#-y
```

shows a simple modifier relationship where x is modified by y. The configurations:

```
x -#---\ /---#-y          x-#---\ /
    |                     |
    |                     |
    |                     |
    #                     #
    |                     |
    z                     y
```

show verb relationships; in the left case above, z acts as the verb relating x and y, in the right case above y acts as the verb and x acts as the subject.

These relationships can be built up into complexes in two ways: where a "node" in a relationship is a loop of dots, the element participating at that node is the entire contents of the loop; where a node terminates in the graphical part of a relationship, the element at that node is the act of making that relationship.

A number of interesting things happen when syntax is "externalized" in this way. Syntax came about originally in conjunction with speech, where speaker and listener are constrained by: (1) the requirement that the listener "decode" the message approximately synchronized in real time with the speaker; and (2) the aid of only whatever "temporary storage" the listener has available in short-term memory. One might say that the function of syntax is to pre-code the message

with storage cues so that the listener will know how to park pieces of the message in short-term memory so that they can be properly assembled in the logical relationships desired by the speaker -- all in more or less real time without getting behind the speaker. Writing, however, changes the picture completely. Obviously, the real-time constraints are absent: the reader may take as much time as desired, may revisit parts of the message as many times as are necessary, and may even browse the message "out of order." In addition, a written document may be said to provide its own storage. In contrast to speech, where whatever parts of the message are not properly stored in short-term memory by the listener are simply (and irretrievably) gone, the written message persists: it stores itself, it stores its structure, it stores its own logical relationships.

Secondly, by externalizing syntax, all points and substructures in the message are accessible in ways not normally found in speech. That they are accessible to the reader has already been discussed. Some interesting ways they are accessible to the writer are revealed by Figure 1. Note the relationship of the phrase "story of the group in isolation" to a larger whole in which it appears. In an externalized graphical syntax, such a relationship is easy to simply draw; joining a part with a larger whole in which it participates is as easy as joining a part with a disjoint part. Relationships between a part and a larger whole in which the part occurs are an obvious logical structure that occurs commonly in the world; yet this is difficult to do in conventional syntax. In addition, the fact that relationships may simply be drawn to parts of the message already laid out allows for complex multiple pathways to be established within even small messages; the message may feed back upon itself. Feedback, while a ubiquitous structure in nature, is notoriously difficult to deal with. It violates the principle set theorists call "well-foundedness"; it may induce the potential for infinite loops in computer programs; where feedback is introduced into the way sound elements are combined in an electronic synthesizer the results may be completely unpredictable: all bets are off. Figure 1 also illustrates this concept of feedback inside the sentence: the "highest-level" logical relationship shown in Figure 1 relates the configuration at the very bottom, in which "denying the volcano" is a modifier, with a cluster "already" deep within the message: "armor : light against the sleep."

A feedback loop may seem an inimical structure to a programmer, where the threat of infinite loop is ever present (and indeed the infinite loop stands out as an archetype "cardinal bug" second only in its fearsomeness to an out-and-out crash); one may say that the threat of infinite loop stands as the fear at the heart of all programming. (Technically, the theorem that one cannot algorithmically determine whether a general computer program will lead to an infinite loop is known as the halting problem, and establishes absolute limits on what is computable.) Yet, when the composer induces feedback into synthesized sound structures, the ear can hear it as a single sound; when a graphical feedback loop is established in a visual syntax, the mind can apprehend the loop as a whole as a single gestalt. Of course to do so, time must not be constrained. It is difficult to see how an aural syntax, subject to real-time constraints, could accommodate feedback loops.

A diagram syntax is notably non-linear. While this is an important point, one must be careful to avoid going too far in pushing non-linearity as a distinction between a diagram syntax and the conventional speech syntax. The essence of syntax is its ability to convey logical relationships across a distance of intervening words; one might say syntax has been our way out of the bind of achieving complex speech structures in the face of the constraint of linear time. Conventional syntax provides a start toward obtaining full non-linearity from an inherently linear channel; a diagram syntax can break free completely to non-linearity without restraint. Non-linearity is freed to extend far down into the fine structure of language -- just barely above the word. Or, to put it slightly differently:

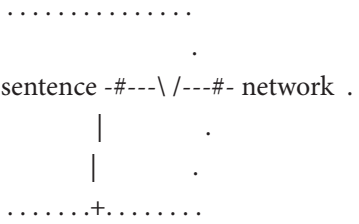

THE INTERACTIVE JUXTAPOSITION

But how to actually achieve juxtaposition of words -- to place them literally on top of one another -- and sacrifice nothing in the way of intelligibility? Too often we think of words simply as whatever comes out of a word processor -- or

perhaps one should call it a word constrainer, forcing as it does the words into the familiar linear chains (with a nod to non-linearity by allowing hypertext links) and certainly not allowing words to be one atop another! A graphics program, on the other hand, allows text objects to be placed on top of one another with complete graphical freedom, but the legibility problem remains. Yet the graphics program gives a clue: juxtaposition combined with intelligibility is achieved (at last) by using interactive software. In a construction I call a simultaneity, words are placed in the same location -- with all the freedom and fluidity a graphics program allows. At first it appears the words are simply overlaying one another -- with no solution at all to the problem of overlay plus legibility. In this state the simultaneity may be called closed. The act of opening the simultaneity consists of moving the cursor using the mouse to a particular "hot spot" on the screen. When the cursor enters this hot spot, all layers of the simultaneity but one are hidden: the one visible layer can be read unimpeded by its partners in the juxtaposition.

Figure 2 shows a simultaneity from Intergrams3. In 2a the simultaneity is closed and all layers are visible; in the detail views 2b-2e the simultaneity is opened showing each layer. (A static illustration cannot convey the tactile aspects of causing the different elements to appear by moving the mouse with one's hand; the reader will have to try to imagine this.)

Figure 2a

Figure 2b

Figure 2c

Figure 2d

Figure 2e

TAKING THE DIAGRAM INTERACTIVE: HYPERTEXT AS A MEDIUM OF THOUGHT

A diagram is a marvellous instrument for presenting information of great complexity in a small space -- to the point that the phrase "Well, you'll have to draw me a diagram" is a stereotype epithet of complaint that something is too

complex. There are limitations to diagrams, however. What happens when the space required is not small? How does one manage a diagram comprising thousands of elements? Enter hypertext4.

Hypertext is most often thought of as a special kind of computer software -- or as the documents produced using that software, but here I would like to consider the idea of hypertext as virtual diagram. In the classical model of hypertext, a document is structured as a network of nodes and links. The nodes are typically either entire documents, or document regions (known as anchors); a link is a relationship between document places such that clicking on the anchor at the source end automatically takes the user to the destination anchor. If a hypertext is small enough and simple enough, the entire network can be represented by other means than using a computer -- on paper, for instance.

Often hypertext begins (alas) at the level of the document; such documents are fully linear and use completely traditional methods for structuring text internally. Using links, associations are built up among places in these documents. The notation of the diagram poems suggests a different possibility: hypertext built up from scratch using very fine-grained word elements, where hypertext is used to carry the infrastructures of language itself, e.g. syntax. One may speak here of hypertext as medium of thought: rather than hypertext serving as an association structure for thoughts that are not themselves hypertexts, an individual thought itself is "entirely" hypertext. To use terminology familiar to computer programmers, hypertext becomes a medium in which one thinks "natively."

Why should we do this: construct a morphemic hypertext5 -- hypertext taken into the fine structure of language? Why not make do with the syntax we have? Why not leave hypertext structure to relate "conventional" documents, at the level known in the hypertext literature as the lexia6?

To answer this question, let me pose a counter-question: How does a single mind apprehend a complex network? It is becoming more and more clear that not only are networks -- in the actual physical sense -- becoming more and more important in our lives, the network as a metaphor is becoming increasingly important in dealing with a wide range of aspects of living. What does it mean for

thought when an individual thought is itself a network? Does it help in understanding the complexities of life's networks around us, containing us, moving us, to "think native" in a mode that is inherently network? Many seek in art a refuge from complexity; indeed, many consider simplicity as such a paramount goal for art that it virtually defines artistic purpose. For others, complexity is taken as a given in this life, and art is seen as an aid that can help us to live with it rather than fight it or withdraw from it. To understand the network one becomes the network. Thought itself is a network, there is no other-than-network:

The obstacles in the way of achieving such a hypertext of thought are many. Among them are:

(1) Lack of Tools. Most commercially available hypertext systems are not adequate. Although much attention has been paid in the hypertext research community to a variety of structural models other than the standard "node-link" hypertext model7, this has borne very little fruit in tools available for the kinds of computers writers are likely to have accessible. Instead, commercially available hypertext software tends to either adhere too rigidly to a node-link model or require the user to build everything "by hand". Typical hypertext structures are or-based, i.e. disjunctive: from lexia L with links X, Y, and Z one may choose X or Y or Z. Syntax structures are and-based, i.e. conjunctive: a sentence with parts X and Y and Z consists of X and Y and Z8. (Consider the classical phrase structure rule:

S -> NP + VP

A sentence can be rewritten as a noun phrase followed by a verb phrase. One does not get to choose which of NP and VP to use; they are both there.) This is not to argue against the use of disjunctive structure, or "classical" hypertext links. Rather, the need is for both to be available as an author requires. Typically, commercially available software has no built-in support for conjunctive abstractions at all.

Another problem with with available software packages is too rigid an attitude toward behavior. Available hypertext systems typically offer only off-the-shelf behaviors that can't be extended by the user. At the other extreme, systems

like Hypercard are fully programmable, but don't allow that programmability to be encapsulated in pluggable objects. (For instance, a Hypercard button has no storage containers!)

(2) Reticence to tackle "language itself." There is no gainsaying that the idea of using hypertext to carry the infrastructure of language itself is an extremely radical proposition -- one from which many will shrink. One source of objection is the idea that "language itself" is off-limits by virtue of being biologically hard-wired9. There are two answers to this: the artistic answer and the engineering answer. For the artistic answer, consider the analogy of dance. No one would dispute that there is a biological basis for how our bodies are put together, for the conformation of bone structure, for the ways that joints work: in short biology places many constraints on how the human body can move. This has not notably abolished the dance. To the contrary: one may say it has created the dance: we admire those who can show us what the boundaries are for how the human body can move, who can take us all the way up to those boundaries and perhaps even stretch them. To the degree that syntax is biological, it makes experimentation on the limits of syntactic structure more valuable rather than less valuable. For the engineering answer, consider the analogy of computer networks. Again: there is no disputing that neurons are biological objects, and that genetics has a great deal to do with how neurons function individually and how the nervous system functions collectively. This does not diminish the utility or importance of those "externalized nervous systems" we call computer networks. The proposal for hypertext as a medium of thought, for hypertext inside the infrastructure of language, is a proposal for an "externalization" of syntax analogous to the externalization of the nervous system manifested in computer networks.10 Just as computer networks do not "replace" the biological nervous system, an externalized mechanism of thought does not "replace" syntax; rather it adds to syntax and allows new possibilities.

For instance: how do we allow more than one user "inside the sentence"? For a diagram syntax this is almost trivially easy: each user's relationships can be distinctively marked -- using color, for instance, or any other form of explicit marking. How is it possible using conventional syntax to construct a "multi-user sentence"? It is exactly in joining multiple users that our biological nervous systems break down and externalized ones show their true value. How does one construct a true multi-user medium of thought? To repeat: a multi-user medium

of thought does not mean a multi-user mechanism for bringing together "single-user thoughts" but rather a medium where the individual thought can be a multi-user construction. Just as multi-user interactions require an externalization of the nervous system, a true multi-user medium of thought will require an externalization of syntax.

It all interacts11:

Notes

1 See Cage's Silence.

2 See Rosenberg's Diagrams Series 3, published on-demand by the author, Grindstone, PA,

1979. Excerpts appeared in Interstate 14, Austin, Texas, 1981.

3 See Rosenberg's Intergrams, Eastgate Systems, Watertown MA, 1993.

4 The term `hypertext' was originally coined by Ted Nelson. The literature on hypertext is extensive; for a bibliography (though slightly dated) see Harpold, Terence, "Hypertext and Hypermedia: A Selected Bibliography", ed. Berk, Emily, and Devlin, Joseph, The Hypertext / Hypermedia Handbook, McGraw-Hill, New York, 1991. The best single-source introduction to hypertext is probably still Nelson, Theodore H., Literary Machines, T.H. Nelson, Swarthmore, PA, 1981.

5 The term `morphemic hypertext' was applied to my work by the hypertext researcher Catherine C. Marshall (private correspondence).

6 The term `lexia' was borrowed from the writings of Barthes by George Landow to refer to a document piece at a hypertext node; see Landow, G. P., Hypertext: The Convergence of Contemporary Critical Theory and Technology, Johns Hopkins University Press, 1992.

7 See for instance Marshall, Catherine C., Halasz, Frank G., Rogers, Russell A., and Janssen, William C. Jr., "Aquanet: a hypertext tool to hold your knowledge in place", Proceedings of Hypertext '91, ACM, New York, 1991 for a model based on relations; Parunak, H. Van Dyke, "Don't Link Me In: Set Based Hypermedia for Taxonomic Reasoning", Proceedings of Hypertext '91, ACM, New York, 1991 for a model based on sets; and Stotts, P. David, and Furuta, Richard, "Petri-net based hypertext: Document structure with browsing semantics", ACM Trans. Off. Inf.Syst., 7, 1, (January), 1989 for a model based on Petri nets.

8 The concept of conjunctive hypertext was introduced in Rosenberg, Jim, "Navigating Nowhere / Hypertext Infrawhere", SIGLINK Newsletter 3, 3, December 1994, http://www.well.com/user/jer/NNHI.html.

9 For a review of issues pertaining to the biological basis of language see Pinker, Steven, The Language Instinct, William Morrow and Company, New York, 1994.

10 Externalization of language is discussed extensively in Donald, Merlin. 1991. Orignis of the Modern Mind. Cambridge: Harvard University Press.

11 The final figure is a single plane in a simultaneity from Rosenberg, Jim, Diffractions through: Thirst weep ransack (frailty) veer tide elegy, Eastgate Systems, Watertown MA, to appear.

THE STRUCTURE OF HYPERTEXT ACTIVITY

INTRODUCTION

A hypertext is a document in which interactive structure operations are intermingled with the text; hypertext structure is usually investigated from the point of view of the "real" structure connecting these operations. E.g. in a classical node-link hypertext, as might be described by the Dexter Hypertext Reference Model,[19] a graph can be constructed on the set of nodes where each edge is identified with a link; structure discussions typically take place with respect to this graph. This overall structure graph may not be apparent to the reader. Readers discover structure through activities provided by the hypertext. This paper will present a framework for discussing the structure of these activities, explicitly based on the reader's point of view. We present a three-layer scheme for discussing hypertext activity: Acteme/Episode/Session. The acteme is an extremely low-level unit of activity, such as following a link. Multiple actemes are combined into an intermediate level unit, which we call the episode1, and at the the high end we will investigate a unit called the session. We will focus much of our discussion on the episode, emergence of the episode from the acteme, the structure of multiple episodes, and how these relate to familiar issues of hypertext rhetoric. The primary focus of this paper will be on literary hypertext, but many of the concepts may be applicable to hypertexts generally.

ACTEMES

This paper takes a broad view of what constitutes hypertext—often narrowly defined as text with embedded links. However, many other structure mod-

19 See Halasz and Schwartz's "The Dexter Hypertext Reference Model."

els have been proposed: sets,[20] relations,[21] Petri nets,[22] etc. We will count as hypertext any kind of system in which text contains embedded interactive structure operations. The lowest level of hypertext activity is to execute such an Seeoperation, e.g. following a link. We coin the term "acteme" to describe this level of activity.

TYPES OF ACTEME

The most familiar form of acteme is link-following—clearly a directional form of acteme. A link may be followed by (1) clicking on an anchor either graphically visible or inferred by the reader; (2) operating an intermediate interactive device showing all possible links, such as a menu of link names; (3) clicking on an overview map (this is really a special case of (1)) and perhaps others. A link menu may contain other information than a link name; in MacWeb [28], link menus contain type information.

There are other actemes pertaining to links. Nearly every form of link-based hypertext allows the user to go back. Hypertext backtracking has been discussed in detail by Bieber. Bieber asks: "Should backtracking trigger an 'undo' operation or simply reflect the current state of the departure nodes?" (PAGE NUMBER). This is an important question, with serious implications for hypertext rhetoric. The simple act of going back may have multiple types. One may revisit a lexia simply to read it again, or it may be a genuine "undo": the reader didn't mean to follow that link at all. These are arguably different actemes, though typically not distinguished by hypertext user interface behavior.

Aquanet[23] uses relations rather than links; for a literary example of relations see Rosenberg's ***Intergrams***. A relation slot is opened or closed; opening a slot is the acteme analogous to following a link. Closing a relation slot somewhat

20 See Parunak's "Don't Link Me In: Set Based Hypermedia for Taxonomic Reasoning."

21 See Marshall et al.:"Aquanet: A Hypertext Tool to Hold your Knowledge in Place."

22 See Stotts and Furuta's "Petri-Net Based Hypertext: Document Structure with Browsing Semantics."

23 See Marshall et al.:"Aquanet: A Hypertext Tool to Hold your Knowledge in Place."

resembles going back for links, but the situation is much more complex. A relation may be n-ary, i.e. may include an arbitrary number of slots. Thus a relation slot may be closed to open other slots, with a clear connotation of continuing rather than backtracking. While a link (even a bidirectional one) is clearly directional, a relation is non-directional, since the relation slots appear in the substructure as peers. There is a mild sense of directionality to relations in the sense that slot contents are "related together", with a directionality inwards from all slots to a central point—one reason why relational substructuring and spatial hypertext are so closely related. (For more on this point see [32], [24], [22], [21].) Also, to use the terminology of [32], links may be described as disjunctive substructuring, in contrast to relations which are conjunctive: whereas one may typically choose any link out from a lexia (from lexia L one may choose link A or link B or link C, etc.) a relation exists among all of its slots (relation R has slot A and slot B and slot C, etc.). Disjunctive substructuring is "or-based" but conjunctive substructuring is "and-based".

Similar to relations, [30] and [31] use a concept called simultaneities, which have unnamed structurally equal slots ([32]); the acteme consists of moving the mouse cursor among different no-click hot-spots, each of which opens a different slot, or moving the mouse cursor out of all of these hot-spots, which closes the simultaneity. VIKI ([23]) includes spatial aggregates (i.e. piles). The acteme here is to click on a partially obscured element of a spatial aggregate, bringing it forward where the whole object is visible. Both spatial aggregates and simultaneities are conjunctive and non-directional.

Storyspace ([6]) offers, in addition to conventional links, spatial placement of "spaces" in a map view; when opened a space may reveal a lexia or a further map. Spaces used in this way resemble piles; the acteme is opening a space. A space can be closed by clicking on an icon from a floating palette. (For a literary example of spaces used this way, see Gess [12].)

Another form of hypertext substructure is the set. HyperSet [29] used an explicit formal set paradigm, and VIKI incorporates sets (collections) as a substructuring method.2 Set-based actemes include choosing a superset (possibly closing the current element) or opening one of the elements of a set. Sets offer a

quite complicated picture. There is a clear notion of "up" and "down" (up to superset, down from set to element) making sets somewhat directional. Choosing a superset is arguably disjunctive; whether opening an element is conjunctive or disjunctive will depend on the specific hypertext.

IS THE ACTEME INDIVISIBLE?

Kathryn Cramer [9] asks the haunting question, "What's inside a link?" Following a link is usually so effortless that it seems nearly automatic. If indeed a link has content then perhaps this calls into question whether link-following should always be considered an acteme. This issue becomes considerably more thorny where links are chosen from a menu. Arguably here the actual acteme is menu-choice, and link-following is a higher-level unit of activity. A menu of possible link names is itself a display of text. A completely open attitude toward text and linking in hypertext would hold that one should be able to make a link to any form of text anywhere text is visible. What if a menu of link names itself contains an anchor? What about links to links? Similar issues have been raised in the past concerning dematerialization of the lexia. (See [26], [32]).

ACTEME "BOUNDARY CASES"

What shall we do, in this analysis, with the lexia? Does linear reading inside the lexia contain / consist of actemes? Should we consider reading a lexia a single unitary acteme? Should perusing the lexia be considered the "null acteme"? (See [32] on the lexia as the "null navigation choice".) Whether the lexia must be linear is controversial ([25] and [32].) A lexia can contain numerous user interface devices (e.g. scroll-bars). Behaviorally, operating a scroll-bar is as complicated as following a link. However, a link is an explicitly structural device in a way that the within-component scroll-bar is not.3 We will leave the issue of within-lexia actemes open. Reading the lexia might be considered a single acteme, or the lexia might be considered devoid of "internal" actemes.

At the opposite extreme, hypertext may be used to carry the very infrastructure of language itself, e.g. syntax. ([30], [31]). In this case the parallel acteme ´ morpheme becomes exact.

THE HYPERTEXT EPISODE

Multiple actemes may combined into a higher-level unit which we call the episode. An episode is simply whatever group of actemes cohere in the reader's mind as a tangible entity. In a node-link hypertext, the episode will probably consist of all or part of a trail or path. Whereas the acteme typically has an identity which is clear from the hypertext's user interface, the identity of the episode may not be so clear. The user may follow a chain of links as part of a process of exploration that may or may not prove fruitful. Simply following a chain of links does not necessarily make these visitations cohere into a tangible entity. The episode is not simply a unit of hypertext history—where any act is necessarily part of some episode; rather, the hypertext experience consists of executing multiple actemes, some collections of which will resolve into episodes, and some of which may not be part of any episode at all. Indeed, part of the hypertext experience may be described as foraging for episodes.

Whether an instance of backtracking is really an "undo" may be rephrased: Does backtracking revoke membership of actemes in an episode? It depends on the circumstances, both of the hypertext and the reader's frame of mind. The reader might revisit a previous lexia to read it again—perhaps for a sheerly "musical" repetition, or to reread a prior lexia based on some resonance or reference in the present lexia. Here one might argue that all of the backtracking history is part of the episode. Or, the reader may be backtracking to undo having arrived at the current lexia by mistake—backtracking to remove from the episode the acteme that caused arrival at the current lexia. The episode is thus a combination of history through the hypertext, the reader's intention, and the reader's impression of what "hangs together". Of course the reader may arrive at a previously read lexia via a different pathway than simple backtracking; in this case most likely arrival at this lexia should be part of the episode.

An episode is obtained from the composition of actemes. For hypertexts using n-ary relations, the nature of the episode can become quite complex. If the slots of a relation do not themselves contain further actemes then the opening of all of the slots of a relation in turn may constitute or be part of one episode. However, more typically a relational slot will contain further actemes, perhaps nested several layers deep. For a highly nested relational structure with a single root relation, shall we say that traversing the entire structure must necessarily constitute a single episode? That seems arbitrary. We may have relations where visitation of the slots in that relation belong to different episodes. Relational structures inherently lend themselves to a hierarchical episode structure.

For set-based hypertexts the episode is likely to consist of both closing elements to open a superset and opening elements of a set. Similarly, where a Storyspace author has chosen to allow the map view to be visible, the episode may be very heterogeneous indeed, consisting of link traversals, space openings and space closings.

IDENTITY/INTEGRITY OF THE EPISODE

Numerous well known issues in hypertext rhetoric can be rephrased as issues pertaining to determining the identity and maintaining the integrity of the episode. Take for instance the infamous "lost in hyperspace" issue. Disorientation in a hypertext may be described as having "lost (irrecoverably) the thread" of the episode. Even though, as Bernstein argues in [3], it may be explicitly part of an author's artistic purpose not to provide a ready "geographic" form of navigation, (indeed, discovery of the geography without "help" from the author may be a deliberate intended effect,) one may still speak of a reader's discomfort in having suddenly lost all episodes. Although even this may be part of the author's intended purpose, most authors are likely to consider it a flaw if "the spell becomes broken". If a reader has "lost track" of an episode but another readily comes to hand, yielding to the new episode and resisting the temptation to "possess" the former episode is one aspect of what the hypertext experience has to offer. One is in some difficulty as a reader, however, if no episode at all comes to hand. The reader is then likely to begin foraging for an episode. Foraging will be a worthy aesthetic experience if (and only if) it succeeds.

Integrity of the episode is also an issue for conjunctive hypertext. Becoming lost in a nested spatially substructured space may not be an issue; the nesting may provide a clear enough orientation that it is virtually impossible to "get (navigationally) lost". Still, the reader here may be subject to a problem parallel to the disjunctive difficulty of losing all episodes: the integrity of the conjunctive episode may fail if the reader is unable to resolve into a single whole "the components of the and"—i.e. if the conjunction fails to come off. This is particularly a danger with heavily nested structures.

In [17], George Landow appeals to the hypertext author not to link the reader into a place where no inviting links will take the reader out. This may be rephrased as the injunction: Don't leave the reader stranded without an episode. A more open-minded approach to the idea of an "episode vortex" would be to appeal to authors to: (1) be conscious of where these places are; (2) be aesthetically comfortable with them; (3) understand how you expect episode foraging experiences to work when the reader hits them; (4) understand how the reader might come out of the episode foraging experience.

IS THE EPISODE A VIRTUAL DOCUMENT?

Various hypertext systems—e.g. MacWeb [27], (see also [13]), have provided for the possibility of virtual documents: documents which are generated "on the fly" by the operation of the hypertext. Here we ask whether we should not consider the episode to be a kind of virtual document. Despite the vast amount of discussion about linkage, there is still a strong temptation to be "lexia-centric" concerning what constitutes "the document". By contrast, consider a common unit of discourse which may span many paragraphs: the argument. Argumentation structure has been a popular topic for hypertext researchers; (e.g. gIBIS [8], Aquanet [20], Sepia [35]). One of the early uses of Aquanet was to interactively implement Toulmin structures [36] to model argumentation. In laying out such an argument, numerous relations are likely to be required; i.e. "the argument" spans many lexia. Likewise in laying out a Sepia Argumentation Space, numerous links are likely to be required. To fully visit an argument is to carry out hypertext activity at least at the level of the episode, and perhaps even beyond to the

session. If operating multiple actemes is necessary to fully visit an argumentation structure, we may truly say meaning is derived through operation of these actemes. I.e. meaning cannot be confined to the within-component Dexter layer (the lexia); meaning derives from hypertext activity in the large. In node-link hypertext, meaning happens through links. (At its most extreme, even the sentence itself can transcend the lexia [32]).

Figure 1 illustrates the episode as a virtual document4. This concept poses some obvious questions: What structures the episode? Should it have a title? Should it have parts? Should the user be able to give it a title? Shouldn't the user be able to save it? (Saving is what we normally do in software with documents we want to keep ...) (User interface questions will be considered in detail below.)

One possibility for how to structure the episode as a virtual document is creation of an explicit gathering interface. [24] illustrates the use of the hypertext system VIKI as a WWW gathering interface. (VIKI is especially suited to this purpose by the richness of its implicit and spatial structuring methods.) Numerous hypertext systems save "global history lists" in which are recorded all lexia traversed; for a true gathering interface this facility needs to be expanded significantly to allow the reader to edit and mark the history for episodes, provide graphical collection of episodes, and so on. At its simplest, such a gathering interface would have commands "begin episode" and "end episode" similar to "record" and "stop" commands common in numerous application program macro recorders. "Playing back" the episode would inject the results into a graphical hypertext browser.5

The availability and characteristics of a gathering interface are directly related to a major question: What is the structure of the episode? Is it in fact linear? With no gathering interface, there is an unfortunate tendency for the episode to linearize—by default—but there is no reason in principle to suppose that the structure of the episode is any less general than the possible structure of hypertext as a whole: the structure of the episode is what the user makes of it given the available tools of the gathering interface. Absent an explicit formal gathering interface, the main tool used in structuring the episode is simply the user's memory.

FROM ACTEME TO EPISODE

In [17] (or see also [18]), George Landow initiated the study of the relationship among episode, acteme and lexia (though not using that terminology) particularly in regard to such questions as: How does the episode emerge from the actemes? How should the lexia and/or acteme be coded for episodes? Landow introduced his "rhetoric of arrival and departure" with reference to the specific acteme of following a hypertext link. Here we generalize these questions to all forms of acteme.

Relational or spatial actemes call for a different terminology than arrival and departure, but the general questions pertaining to episode/acteme/lexia remain. Consider n-ary relations. Does opening each of the slots of a relation in turn belong to the same episode? Where relational structures are nested, the logical way to read them might be by depth-first traversal. The degree to which different slots of the same relation would belong to different episodes would tend to depend on the complexity of the structure of the slots. If each slot is a lexia with no internal structure, opening all of the slots might naturally fall into one episode. If two slots each have highly complex nested substructure they might fall into separate episodes. Should a hypertext be coded for this difference? When an acteme involves opening a space, should the reader be given a clear graphical cue as to how complex that space is?6 This may be implemented using icons which are miniature graphics of their entire nested structure. (See Figure 2.) Such a miniature prepares the reader for what will happen when the slot is opened.

For spatial hypertext, spatial proximity is one way actemes may be coded for grouping into an episode—if this use of spatial proximity is not preempted by some other structural purpose. Regions containing spatial actemes can be enclosed in a graphic device, such as VIKI's collection frame. Where a node-link hypertext offers a graphical view, links can be coded for episode, either graphically, through names, or both. However, where hybrid methods are used the situation is much more complicated. Consider a Storyspace hypertext in which the "map" view is enabled. Coding for an episode that visits multiple spaces is now much more difficult. Likewise, in set based hypertext the episode is likely to consist of

multiple transitions between "up" navigation and "down" navigation. How should an author code for episode in this case?

Given that the episode is really the reader's structure, to what extent should the writer code for episode at all? This is an aesthetic issue, likely to yield a variety of points of view. (And even a single writer in the context of one work might choose to vary the extent to which episodes are coded.)

MULTI-EPISODE STRUCTURE

In most cases, reading a hypertext will result in several episodes. We now pursue the question of what type of structure might relate some of these episodes.

Multi-Episode Structure is Emergent

The episode itself emerges from reading activity; although the writer may employ coding devices or hints to guide the reader in forming an episode, ultimately the episode is more the reader's structure than the writer's. Structure among multiple episodes is even more likely to be emergent, and is most likely not specifically embodied in formal structural devices of the kind articulated in the Dexter Hypertext Reference Model. Hypertext systems for expressing emergent structure have been studied in detail by Marshall and her colleagues. ([22], [23], [24].) They have found consistently that where a gatherer is unsure about final structure, spatial methods tend to be preferred—often even when more formal structures, such as relations, are available. Even if the formal structure underlying a hypertext uses the familiar node-link model, and is completely disjunctive (as in the section "Types of Acteme" above), where multiple episodes are being gathered using (say) spatial methods, the structure that results from this gathering may be conjunctive rather than disjunctive, or may be a complex combination of conjunctive and disjunctive substructures. I.e.: Even assuming that purely disjunctive methods are sufficient for the author's purpose, the availability of a richer structure palette—specifically including conjunctive substructuring method—may be of serious benefit to the reader. The appeal for a gathering interface issued above must be reissued in even stronger terms as we consider how the reader is to work out the structure among episodes.

CONTOUR AND GAP: THE GEOGRAPHY OF EPISODE

Michael Joyce has written frequently about hypertext contours ([4], [16]). Contour is a multifaceted concept ([33]). Some aspects of contour may be rephrased as questions: How does the reader perceive the episode density of the lexia? How does the reader associate multiple episodes with a map view of the hypertext? How does the reader locate lexia which are particularly rich "hinge points" joining multiple episodes? Fascinatingly, in [15], Terence Harpold investigates what may be described as the "skew-symmetrically opposite" concept. He describes a model of a hypertext as knotted threads; in a kind of counterpoint to Joyce's contour, he describes a concept of gap which we might paraphrase as a void around which episodes may bend but into which no episode reaches. Both concepts concern the geography of episode: in the case of contour, where the episodes are; in the case of gap, where they are not. Where a node-link hypertext has a strong geographical map view interface, the episode yields a trace on this map. Visually associating multiple such traces is an obvious method of structuring multiple episodes.

There may be no map view at all in the hypertext system in which a document is read. Should the reader then create such a map—at least conceptually? Should a gathering interface provide a map display for the reader's gathering activities, even if the "original" hypertext system in question doesn't? Where a map display is available, it is likely to show the writer's "inherent" underlying structure, e.g. for node-link hypertext the node-link map itself.7 How is the trace of episode to be made visible on this map? Bread crumbs ([2]) are a standard device for exhibiting hypertext history on a map view (MacWeb does this, for instance)—but history and episode are two different things altogether. Clearly the reader could use some help here.

What of the geography of episode for non-link hypertexts? Both sets and relations may give the reader a three-dimensional feeling: as a set or relation slot is opened, the region of hypertext thereby made visible may appear — conceptually at least — in a different plane behind or in front of the plane where the user was. The episode may be a kind of tube that transcends multiple planes — pos-

sibly zigzagging "forward" and "backward" (or "up" and "down" in the case of set navigation) multiple times. As in the case of the node-link model, episodes may intersect. A particular lexia may be very rich in episodes, or they may be planes that are strangely bypassed. For relational hypertext where the relation slots are visited in separate episodes, the relation structure itself may be said to relate these episodes; Harpold's metaphor of hypertext as knots seems particularly apt — though the knots may be nested, like layers of an onion.

NARRATION—A LOGIC STRUCTURE OF EPISODE

Narration is an immense issue; a discussion of narration as a whole is beyond the scope of this paper. We only note a few issues here. In attempting to put a narrative structure to a hypertext, surely the reader is attempting to relate not just lexia, but episodes as well. Indeed, the whole concept that a sequence of hypertext activities works together as a single story fragment may be one of the ways by which the reader constructs a concept of episode in the first place. (See [10] for examples.) Whereas above we were concerned with the geographical relationship among episodes, here we are concerned with a logical relationship. (These may or may not be the same.) An emerging logical or narrative schema may have a great deal to do with how the reader forages for episode; as Douglas observes, an emerging narrative picture may have gaps; it is precisely to fill those gaps that the reader may forage for more episodes. (And as Harpold observes, not only is there no guarantee of success, the writer may intentionally make it impossible to find such a "missing key".)

USER INTERFACE ISSUES PERTAINING TO THE EPISODE

The user interface implications of the concept of episode range from very small-scale (e.g. minor details concerning how bread crumbs should work) all the way to a full-scale gathering interface. Consider bread crumbs. A typical bread crumb device shows only history. It is typically oblivious to the question raised above of whether backtracking is an undo; it treats every lexia visited equally, whether the reader is in the midst of a very intense episode or has lost the thread completely and is foraging for a new episode. Clearly it would be useful for bread

crumbs to be typed. Just as MacWeb allows links to be typed, the reader may need typing as well: visitations may need to be typed. Visually this could be indicated in many ways: color coding, icons, etc. A related issue: does the reader need the ability to name the episode? Surely in some cases this would be useful. For named episodes, a bread crumb could be a fully clickable icon which would expand to show (or accept) the episode name.

The go back command found in almost every hypertext system should allow for qualification: Whether backtracking is an "undo" or not must be answered by the user! Likewise, the user must determine whether backtracking should or shouldn't be recorded in the trace of the episode.

The typical save command needs considerable enhancement. In most hypertext systems, the only things which can be saved are: (1) References to particular lexia (bookmarks); (2) the state of the entire hypertext session. We argued above that the episode may function as a virtual document; if so then the reader should have the ability to save it. It is ironic in the extreme that despite all the emphasis on linkage over decades of hypertext research, it is the lexia which is typically saved, not the linkage! The ability to save an episode provides an opportunity to name it, of course.

More elaborately, gathering must be regarded as an important aspect of the hypertext act. Alas, the requirements of a full gathering interface are considerable.

• A gathering interface must provide a rich palette of structuring methods, specifically including spatial structuring methods such as those implemented in VIKI.

• A gathering interface must be at least partially automatic. It should have facilities similar to common macro recorders, so that when the user has indicated that an episode should be started, further activity is automatically assigned to the episode without the user having to do so manually.

• The history mechanism should be available for retroactive editing allowing an episode to be reconstructed after the fact. This is important: just as the episode is emergent structurally and spatially, it is emergent in time as well: you may not realize you are in the midst of an episode until well after it has already begun. The reader must thus be able to edit the history list and gather into an episode actemes already performed.

• A gathering interface is explicitly one hypertext system operating on another; ideally the authors of the hypertext systems at both ends of this transaction would be sensitive to the needs of being plugged into a companion. The writer's hypertext system should have sufficient hooks that an off-the-shelf third party gathering interface can be plugged into it; the reader's gathering system should use sufficiently general system mechanisms as to allow for operation of a variety of hypertext systems.

THE HYPERTEXT SESSION

There is a clear break in hypertext activity when the user quits. An excellent discussion of issues pertaining to the hypertext session may be found in [10]. Douglas's main focus is the issue of closure: how does reading a hypertext "come to an end?" There are all manner of reasons why the hypertext session may end. We examine some of these.8

(1) The session may end due to accident or external circumstance. Perhaps the phone rings, or the power fails, or the computer crashes. It is tempting to simply dismiss this as a not very interesting null case, but it is precisely by measuring the sense of loss at an artificial termination that we may properly assess what needs to be saved from the session. How does the reader recover not only the lexia but the episode as well? Can the episode be recovered? If the reader is associating multiple episodes, can that be recovered? Should it be recovered?

(2) The reader may simply give up after a fruitless search for episode. Similarly the reader may suffer sheer episode fatigue: episodes are at hand, but they seem so similar to episodes already undertaken that the reader simply quits for want of "something fresh".

(3) By contrast, the reader may have achieved a complete sense of episode satiation. This is not necessarily the same concept as closure, as discussed by Douglas. Particularly in a large poetic work, the reader may have no sense of completion in a logical or narrative sense, but may be satiated in a purely imagistic way that makes it seem fruitful to put the work aside for a time. There are some interesting aesthetic issues here. If I have reached episode satiation, I might not want to resume in a subsequent session exactly where I left off, but might instead want to forage "as far away as possible". (Returning in a subsequent session to the same neighborhood where I left off might actually be overtly disappointing.) How do I as a reader do this?

(4) The reader may have reached a tangible "success point" in gathering. (This is probably the closest concept to Douglas's description of closure.) Of course, the reader may not have a formal gathering interface; the gathering in question may simply be formation of a mental map. To borrow Michael Joyce's topographical bent: the reader may quit because of a feeling of having reached a point on the landscape from which the vista seems complete. Or as Douglas puts it, the reader is satisfied that enough logical questions are answered that there is no need to continue. With the luxury of a formal gathering interface, the reader may obtain a sense of completion about the gatherings; i.e. the reader's sense of completion is exactly a writer's sense of completion: the gathered result "works" artistically as-is, now is a good time to stop.

CONCURRENCY OF EPISODE

Some hypertext systems are explicitly designed for concurrent operation by multiple users. (See e.g. [8], [35], [11].) The study of actual concurrency in literary hypertext is surely a worthy subject of research, but here we will investigate the metaphor of concurrency of episode in the context of activity by a single user. In [5], Bernstein makes the intriguing proposal that we personify episodes, endowing a hypertext with what he describes as characters. Characters, of course, exist in a narrative space concurrently (as do Bernstein's). Does it help the reader to imagine episodes as occurring concurrently, even if they are not experienced that way? While this might seem to be stretching a point, consider that as gather-

er, the reader may be assembling a new hypertext containing the gathered results. In this case we have not only the episodes in the original hypertext, there may be potential episodes in the gathered product. Concurrency here is quite literally real in that: (1) potential episodes in the gathered result exist concurrently with the episode in the "original" hypertext; (2) a formal gathering interface may make it exceptionally easy to keep multiple potential episodes "open" at once.

THE READER-AS-WRITER'S ACTIVITY STRUCTURE

It is a commonplace in hypertext rhetoric that the reader is also concurrently a writer; we now explore this from the standpoint of activity. The activity of following a link (the reader's link!) we classified above as an acteme — a low-level unit of hypertext activity. What of a link created by the reader? Shall we describe creation of a link as an acteme? Creation of a link might consist of: (1) selecting the text in the source lexia to serve as the source anchor; (2) telling the system we want to create a link; (3) navigating to the target lexia; (4) selecting the text to serve as a target anchor; (5) informing the system we are completing the link; (6) choosing a name for the link; (7) perhaps choosing a type for the link. This is hardly a low-level unit of activity! The supposed symmetry between reading and writing, from the standpoint of activity structure, is completely illusory in the sense that it may take nearly an order of magnitude more effort to create as a writer what the reader experiences as a simple acteme. What does it take for the writer to create an episode? What has happened to the reader's episode while the reader-as-writer is creating a link? This is a most unpleasant question! Is the episode "in suspension"? Is creation of the link simply a part of the episode? What is the risk that creation of the link will "break the spell" of the episode?

In place of reader-as-writer, consider the concept of reader-as-gatherer (given a formal gathering interface). A gathering interface, particularly one that implements spatial methods such as VIKI, may serve as a much more lightweight interface than a full-scale authoring environment. For instance, adding a lexia to a pile already open in a gathering interface is likely to involve no more activity than dragging the lexia onto the pile; with a gathering interface designed as such it may even be simpler. A true symmetry in complexity between the reader's acteme and the gatherer's acteme may in fact be achievable. Given that creation of

a simple link may be so much more laborious than following one, one is tempted to ask how much hypertext gathering is actively discouraged by asking the reader to use a full heavy-weight hypertext authoring environment instead of lighter-weight gathering tools.

1Bolter [7] used the term 'episode' the way the term 'lexia' [19] is now conventionally used, but his use of the term episode did not catch on; apologies if this new use of the term episode causes confusion. It was tempting to borrow Douglas's [10] term 'strand' — but this seems to prejudge the issue of whether the episode is linear and to be more narrowly suited to the specifics of the node-link model.

2This author knows of no literary hypertexts explicitly based on sets as a substructuring method. The complete absence of set-based literary hypertext is both striking and hard to explain (though set-based substructuring is not usually present as an off-the-shelf abstraction in commercially available hypertext software.)

3Trellis ([34], [11]) provides a formal basis for dealing with such questions. For a Trellis hypertext one may describe an acteme as any form of hypertext activity which causes the Petri net to fire. If within-component scroll-bars are devices maintained entirely by a client which does not fire the net when they are operated, they would not be considered actemes.

4Figure 1 as drawn implies that the lexia is "atomic" with respect to episodes — i.e. a lexia is either entirely in or entirely out of an episode. Of course an episode may include only part of a lexia; there is no guarantee the reader will read the whole thing.

5Zellweger [37] discusses implementation of a similar concept, though her paths are constructed by the author rather than the reader.

6At the Spatial Metaphors Workshop at ECHT'94, Mark Bernstein raised the question of how the user of a hypertext might be able to estimate the cost of following a link. Should actemes be coded so that the reader can estimate the cost of activating them?

7In [10] Douglas refers to such maps as "cognitive maps". To call the "actual" map of the node-link structure a "cognitive" map is a serious confusion. The map may be structural more than it is cognitive. One might in some cases call a

reader's map in the context of an overt gathering interface cognitive; whether the writer's structure map is cognitive or not depends on the circumstances.

8This section is heavily indebted to Douglas's paper. While it should not be taken as simply a restatement of her work, most of the ideas in this section were the direct result of reading her very stimulating discussion.

AND AND: CONJUNCTIVE HYPERTEXT AND THE STRUCTURE ACTEME JUNCTURE

A rhizome has no beginning or end; it is always in the middle, between things, interbeing, intermezzo. The tree is a filiation, but the rhizome is alliance, uniquely alliance. The tree imposes the verb "to be," but the fabric of the rhizome is the conjunction, "and . . . and . . . and . . ." This conjunction carries enough force to shake and uproot the verb "to be."

Deleuze and Guattari [11, p. 25]

INTRODUCTION

The concept of conjunctive hypertext was originally introduced in [38] (see also [39]). Normally we associate node-link hypertext with the ability to choose links: if a given lexia has links A, B, C the user can choose A or B or C. Perhaps no one else has put this association of hypertext with choice in quite so succinctly lyrical a way as Shelley Jackson: "Hypertext is the banished body. Its compositional principle is desire."[24] However, there is another possible relationship between a whole and its parts: a construct may have components A, B, C, in which the construct consists of A and B and C. Because in logic the 'and' operation is called a conjunction and the 'or' operation is called a disjunction, the term conjunctive hypertext refers to hypertext constructions where the relationship between a component and its elements is 'and' rather than 'or'. I.e. whereas disjunctive hypertext presents activities as alternatives, conjunctive hypertext presents activities as elements to be combined into a whole effect. While presentation issues are important in a discussion of conjunctive hypertext, conjunctivity does not derive simply from a presentation method. Rather, conjunctivity is a particular attitude toward how multiple activities in a hypertext construct relate to one another. Conjunctive constructs arise quite naturally in spatial hypertext, which has been a very active area of study in recent years. Thus it seems opportune to review issues pertaining to conjunctive hypertext generally. In pursu-

24 I am indebted to N. Katherine Hayles for this quotation. See "Flickering Connectivities in Shelley Jackson's *Patchwork Girl*: The Importance of Media-Specific Analysis."

ing conjunctive vs. disjunctive hypertext, we are inquiring into the relationship among actemes and the relationship of actemes to structure. (The terminology of [39] will be used throughout this paper. The term acteme refers to a very low-level unit of activity, such as following a link. The term is useful as a generalized unit of activity that can be applied in situations where there may be no links.)

It is important to emphasize at the outset that this paper is in no way maksing any kind of claim whatsoever that conjunctive methods are somehow "superior" to disjunctive methods. Authors should have available the widest possible variety of methodologies, and they should all be available together.

THE MULTI-ACTEME CONSTRUCT

The most familiar hypertext structure is a simple binary link.[25] The link is associated with an anchor at each end; although the link may be bidirectional, unidirectional links are more common. In this case the link is activated by (typically) clicking on a point within the bounding region of the source anchor. The focus then changes to the destination anchor, which most often involves either changing the user's location within the same lexia or opening a new lexia. Thus the binary link is a structure with a single acteme: following the link. This section considers constructs that contain multiple actemes. Of course a multi-acteme construct need not be conjunctive; indeed the relationship among actemes in a multi-acteme construct may even be ambiguous.

An example will serve to illustrate this concept. Figure 1 shows a construct I call a simultaneity taken from [40]. The aggregate entity consists of a spatially overlaid cluster, whose members are either phrases or other clusters. As the mouse approaches the cluster, it "opens", revealing one of the members of the cluster. At the same time a stacked set of frames appear. These serve as "on mouseOver" style hot-spots; when the mouse enters one of these hot-spots, the layer associated with that member of the cluster comes to the top and is displayed.

25 By "binary link" is meant a link with a single source anchor and a single destination anchor; it is a binary structure in the sense that the two endpoints are joined by the link. It could just as easily be described as a unary structure, by analogy with unary functions, in that following a link yields a single destination anchor. Referring to a typical link as binary rather than unary seems less confusing when considering the possibility of multi-headed and/or multi-tailed links.

Thus the cluster has an acteme for entering each member, which all act as peers. The cluster consists of all of its members; the members do not function as "alternatives" any more than the various parts of a sentence act as alternatives to one another.

Note that the relationship of the actemes in a construct like the simultaneity described above is quite different from the relationship of actemes composed into an episode (in the terminology of [39]). Composition of actemes results in a larger-order unit, such as a trail of links, or in the conjunctive case perhaps descending into a hierarchy and visiting an entire sub-hierarchy. Typically the episode need not be localized; in speaking here of a muti-acteme construct we are speaking about a single localized construct.

ACTEME RELATIONSHIPS MAY BE AMBIGUOUS

In speaking about a mult-acteme construct, the term "construct" may or may not mean "structure". A system such as VIKI [33], CAOS [37], or VKB [43], may allow spatial relationships that are by design ambiguous and whose structural relationship is unclear. Consider a "brainstorming session" where you aren't really sure whether an idea is related to those you've placed in a named collection. You might place it "close to" the collection. By doing so you've placed it where you won't forget about it when you consider the collection, but it isn't in the collection because you haven't decided that that's where it belongs. "Close by" is a type of relationship one might describe as parastructural. Artistic and literary works often present examples of parastructure. For instance, passage [7] by Philippe Bootz contains actemes where the presence of the acteme is not even clear; if you don't activate it within a certain window of time you may have lost the ability to activate it at all.

Sets provide another arena where structural relationships may be ambiguous. (The pioneering reference on set-based (taxonomic) hypertext is [36]. Taxonomic hypertext remains at this writing an underutilized form of hypertext, particularly in literary hypertext.) Does a set consist of "all" of its members? For a given element, what supersets does it belong to? Is one supposed to "choose" a superset? Does one have to "visit" all of them? Answers to such questions may be a matter of context. (Structure completion is discussed in more detail below.)

A more complex case is provided by the n-ary link — a link where there may be an arbitrary number of associated anchors. Several systems support n-ary links, such as Chimera [1], HyperDisco [45], Webvise [16, 18], and the HyTime aggregate link [14]; the n-ary link is explicitly supported in the Open Hypermedia Protocol (OHP) [10]. Unfortunately, the literature on user experience with n-ary links is not extensive. In a typical case an n-ary link will have a single source anchor and multiple destination anchors. "Following" the link means opening all of the destination anchors — as opposed to opening a menu which would allow the user to choose which destination anchor to open. In this sense the n-ary link is a conjunctive construct, yet if there is only one source anchor there is only one acteme, so it would appear that we don't have a multi-acteme construct here. However, once all the destination anchors are open, each destination lexia may contain actemes of its own. At a minimum, each destination lexia is likely to have a close acteme. What is the relationship among these actemes? If they can all be opened at once, can they all be closed at once? [41] discusses the concept of non-unitary location. "Where" is the user after an n-ary link has been activated? Once destination lexia have been activated, what is the relationship among the original set of destination anchors? Is it even clear?

Of course, everyone is familiar with at least one form of potentially multi-acteme construct: the lexia itself. The structure of the lexia is less well studied than it should be. Often the lexia is simply considered to be an atom in hypertextual structure; the real meat of hypertext is thus presumed to be elsewhere. A lexia which contains several links might be presumed to structure those links by the content of the lexia. This presents actemes whose relationship to one another is as general as language itself.

PEER RELATIONSHIPS

The simultaneity structure discussed above is an example of actemes in a peer relationship to one another. Peer relationships may have the following properties:

- The construct as a whole has behavior.
- A navigation interface allows movement from one peer to another.

• An operation may affect all components at once. E.g. if the construct has a close operation, when the construct is closed, all of the component peers are closed.

Another example of a peer structure is the tabbed dialog box familiar in many operating system user interfaces. (However a tabbed dialog box is most likely not conjunctive, since the user can simply choose which tab to view in the same way as choosing links.)

IMPLICIT SPATIAL RELATIONSHIPS

An extremely useful form of implicit spatial relationship is the pile — objects placed "loosely" in spatial proximity. Piles have been supported by spatial hypertext systems such as VIKI, CAOS, VKB, and Web Squirrel [2]. An interesting study of the pile as a user interface metaphor was given by Mander et al [29]. Their prototype included some useful concepts, such as the ability to spread out a pile by means of a mouse gesture. Current spatial hypertext systems are somewhat primitive in their interface provisions for pile handling; they typically require an object to be selected to bring it to the top. This poses serious problems for an object which becomes completely occluded. (And it could be argued this makes the conjunctivity of the pile construct problematic.) An interface that treated a pile as a peer structure would be extremely useful.

A pile is often not permanently represented by an internal structure in spatial hypertext systems, but is computed on the fly by a spatial parser. E.g. a double-click may select an entire pile, yet not present that pile to the user as a specific structure in a hierarchical display of a document's objects. Is the pile a structure, or a particular artifact of behavior?[26]

Another kind of implicit spatial relationship is the grid: elements arrayed in vertical and/or horizontal alignment [42]. The conjunctivity of grids is especially clear, since they are normally presented so that all elements are visible.

---

26 Integration of behavior into the structural computing paradigm emerged as an open question at the Second Workshop on Structural Computing, San Antonio 2000.

TIME RELATIONSHIPS

Instead of being associated spatially, elements of a hypertext may occur in the same space associated by time. For instance, two elements might oscillate back and forth. (For an interesting and undeservedly obscure study of some of the larger issues raised by oscillation in hypertext, see [21].) More generally, a sequence of elements may be presented in a loop. In this case conjunctivity arises from the elements "all being there" through time rather than space.[27]

Formal Relationships

There are numerous ways in which multiple elements may be presented in a formal relationship. The "most formal" of these is the relation, as implemented for example in Aquanet [31].[28] A more familiar form of formal relationship is the hierarchy. At its simplest, the hierarchy consists of a single parent object and some number of contained objects; in hypertext theory this is usually referred to as a composite [19, 15, 17 chap. 7].

ACTUALIZING THE CONJUNCTION

"Actualizing" a binary link is a simple matter: the user clicks on the link and is taken to the target. For a conjunctive construct with possibly many actemes, "putting together" the terms of the conjunction raises several much more complicated issues, which are dealt with in this section.

CO-PRESENTATION

27 The familiar animated gif is certainly a loop, but has no actemes.

28 In my opinion, the "lesson" of Aquanet has been widely misread. User experience research with Aquanet determined that users are often unwilling to commit to "hard" structures in advance (see Marshall and Rogers' "Two Years Before the Mist: Experiences with Aquanet"). This led directly to the creation of VIKI, with extensive support for emergent and ambiguous structure. This support is extremely welcome. However, the Aquanet experience has been read as evidence that users do not want to commit to "hard" structures ever, at all — and thus support for relational structures has disappeared from hypertext systems. This is unfortunate. Explicit structuring such as relations, particularly with Aquanet-like graphical rendering, should be provided along with the more implicit and ambiguous forms of structure in VIKI, CAOS, and VKB. Since we have no hypertext systems in which both forms of structuring are available, there is no data on the basis of which one might conclude — as the results from Aqanet are often portrayed — that when both forms of structure are available that users simply avoid hard explicit structuring completely.

The easiest method of actualizing a conjunction is to simply present all the elements on the screen at once. The most common instance of this is the HTML frame set.[29] Co-presentation has some obvious difficulties. Screen real estate limits how many elements can be presented together. This can be particularly awkward if the elements are themselves compounds. A concept like the HTML frame set works well when there is only a single top-level structure (the page) and a single set of frames; with nested structures the navigation interface can become confusing. The pile represents a particular challenge for co-presentation, as already discussed.

The concept of tabletop [44, 30] provides an interesting antecedent to co-presentation, though with a completely different purpose. Tabletops allowed co-presentation of multiple lexia in a hypertext in a pre-programmed way as a kind of illustration of a juncture point in forming an episode. (As used, tabletops tended to be disjunctive, like the hypertexts they typically commented upon.[30])

For an interesting literary example of co-presentation used in a straightforward way see [8], which is implemented using HTML frames.

DELEGATED PRESENTATION

A layer of hypertext infrastructure may delegate to some other layer of software the actualizing of a conjunction. E.g. an Open Hypermedia Systems (OHS) component may link-enable productivity applications such as a word processor or spreadsheet. If an OHS link service responds to clicking on an anchor with a multi-tailed link, it may be left to the native operating system windowing environment to display the destination lexia, where it is unlikely that there is any real "awareness" of the conjunctive nature of the n-ary link. Once the link service has completed its work, any awareness that the anchors retrieved are in any sense peers is likely to be absent. Needless to say, this is a fairly weak form of

29 David Stotts considers the HTML frame concept to be the most common example of a multi-tailed link [personal communication]; but it seems more reasonable to view the entire frame set as a link target—a composite perhaps.

30 The self-referential character of hypertexts with tabletops should be of great interest to the literary community, considering the widespread use of self-referentiality in postmodern literature.

conjunctivity. If a document author requires presentation behavior for actualizing a conjunction that embodies conjunctive awareness, this may make delegated presentation problematic.

TRAVERSE AND THE CONCEPT OF FOCUS

Our notions of both of these two concepts have been mainly formulated with respect to the binary link. Thus we tend to think of the focus being "at" a single node; traverse takes us to another single node. [41] discussed the concept of non-unitary location: in conjunctive hypertext one may located "in" many places at once.

PEER TRAVERSE

Conjunctive structures have a special kind of traverse: traverse among peers. Discussions of composites typically ignore this issue, and many user interfaces do not allow easy navigation among peers. Instead, one often has to return to a parent structure to navigate among peers. This yields a particular structural form of cycle, in the sense of [4]. (Bernstein's concept of cycle lumps together a number of disparate concepts, which require separate elaboration. Several will be examined in this paper.) This type of purely "mechanical" cycle — which could be eliminated with proper peer navigation — is quite different from a genuine case of "revisitation" along a trail of binary links. Peer-to-peer traverse is important in conjunctive hypertext to "put together" the conjunction in cases where co-presentation is not practical. Just as it might take multiple readings to resolve a sentence containing a difficult idea, it might take multiple peer traversals to resolve a conjunction. (This might also give the appearance of what Bernstein would describe as a cycle.) It is interesting to note that Grønbæk and Trigg cite traversal as a key element in distinguishing composites from links: "Containment and opening are suggestive of a composite, while connection and traversing are suggestive of a (possibly multiheaded) link." [17, p. 87] This distinction completely misses the issue of peer-to-peer traverse.

An interesting question regarding peer traverse is where this functionality should reside. Advocates of composites might argue that no "special" behavior is required here; if an author wants peer traverse, that can always be arranged by

specifically creating links among the peers. This can raise difficulties, particularly in an OHS context. Perhaps the conjunctive structure is aggregating elements over which one has no authorial control. In this case the peer navigation needs to reside either in an external structure base or with the higher-level construct. Whatever structure server is providing the terms of the conjunction needs to provide the peer navigation mechanism as well. Indeed a lexia which is one among several peers in a conjunction may not "know" it is a peer, or what its other peers are; this depends on what conjunctive structure has been activated. It would certainly be possible for a peer lexia to contact a link service to inform it of other peers (or to inform it of a higher-level "controller" object), though of course that would require the link service to maintain a high degree of state. Even given a link service which supported peer or controller identification and navigation, the service presenting a peer would have to "know" it needed to contact the link service to obtain this information. In the example presented earlier in Figure 1, peer navigation was of course built in to the interface.

SUBSCREENING

Another method of dealing with traverse in conjunctive structures may be described as subscreening: An area of the screen is divided into sub-regions, with traverse allowed independently in each sub-region. This is the approach taken in HTML frame sets. This has numerous potential complications for the user interface. There are now two concepts of "back": back within the subscreen, and back for the higher-level unit as a whole. It is quite common with HTML frame sets that there is no visual cue that subscreening is in effect: you often don't know you've been "captured" inside a frame until you follow a link and find that traverse only occurred inside a subscreen. When subscreening is occurring, what should be bookmarked? The discomfort one may have at answering this question is one of the clues that location in conjunctive structures is not always unitary.

Subscreening is subject to a particular form of disorder that may be called anti-conjunctive drift: the conjunctive relationship among subscreens may at first have been clear, but as traverse occurs in each subscreen, this relationship can deteriorate. HTML frame sets illustrate an extreme version of this problem. Frames are sometimes used to provide links but still "capture" the reader "within" a particular URL. While normally we would think of a conjunctive relationship

as "cooperation" among the terms of the conjunction, in this case it can deteriorate to almost a kind of parasitism. Anti-conjunctive drift works quite counter to peer-to-peer traverse; indeed one could describe this problem as a sudden feeling by the reader that objects that are supposed to be peers no longer are peers. Interestingly, there is no common interface which under these circumstances will "release" the association of subscreens — e.g. to make a separate window out of each subscreen. (Though of course some web browsers allow opening a single subscreen as a new window.) Release of subscreens might help considerably to counteract this problem. It would represent a recognition in the user interface that what had been the terms of a conjunction are no longer operating that way and should henceforth operate disjunctively.

PENDING STRUCTURE

When the familiar binary link is followed, nothing in this structure is left pending: the user leaves the source lexia, arrives at the destination lexia, and the transaction is complete, so to speak. Conjunctive structures (and n-ary structures in general) raise a very large issue: as the conjunction is being actualized, this actualization may be incomplete. Some method must be found to indicate to the user what parts of the structure are still pending. Indication of pending structure is an ancient problem: syntax itself may be described as a pre-coding of the message so that the mind knows where to park pieces of the text for processing later to pick up pending parts of the structure.

Does a structure "have to be" completed? Structure completion is related to the literary concept of closure; where closure deals with a feeling by the reader that experience with an entire work has reached a kind of completion, structure completion may be described as the issue of closure in the small. (For a discussion of hypertext closure, see [12].) Scale does matter here: discussions of closure for a hypertext work as a whole are complicated by the problem that the number of potential paths through the whole work may be overwhelming, but the number of "slots" in a structure is typically small. Many hypertexts — particularly literary ones — do not provide a systematic display of "all" of the lexia, but conjunctive structure is normally displayed so that all slots may be visited — whether by co-presentation, peer navigation, or some other method. However this picture becomes less clear when structure slots contain complex structures which may

require a descent of several levels for completion. In this case structure completion begins to resemble closure as an issue.

An interesting question is the issue of whether the notion of conjunction "requires" completion. After all, it might be argued, if structure X consists of all of A, B, C, ... how can we say X has been visited unless all of the components have been visited? This is a matter of context, and the intentions of the hypertext author. We do tend to believe that a conjunctive structure such as a sentence should be read in its entirety or not at all; it would be very strange to say of a document, "Yes, I read a good deal of it, but I just picked the part of each sentence that looked like the best." On the other hand, we often are selective in not reading "all" of a work without requiring that the components of a work be thought of as alternatives. One can imagine a dynamically computed conjunctive structure in which elements continually arrive and depart; the elements are not alternatives to one another but visiting all of them is impossible (by design) because in the time it would take to visit all elements some of them have disappeared and others have been created. Conjunctive structures may be described as somewhat biased toward completion "in the small", but as conjunctivity scales up to large structures, there is little difference on this issue between conjunctive structures and disjunctive structures.

How is pending structure to be indicated? Bread crumbs [3] are a familiar device for indicating structure already visited. One approach to the display of pending structure would be to use such a graphic marking device to display those parts of the structure already visited, with the presumption that the unmarked parts are pending. Current spatial hypertext systems are somewhat weak in their facilities for displaying pending structure. It is typically assumed in the case of a pile, for instance, that all members are visible enough to be accessible by clicking. If a pile member becomes completely occluded, there is typically no method for indicating that it is still pending in visiting the structure: spatial hypertext systems tend to rely completely on co-presentation.

Many other methods of indicating pending structure are possible. For instance, a system might use some form of graphical biasing so that pending mem-

bers of a structure are brought to the top of a display or indicated more prominently in some way.

The notion of pending structure is related to some classical issues in hypertext rhetoric. George Landow has argued that in spite of the extent to which following a link may seem to be an atomic transaction, the way should be eased at each end of the link by a rhetoric of departure and arrival [26]. Where structure is left pending, this rhetoric must be considerably extended. To departure and arrival must be added "pushing" and "popping" — as well as perhaps "rotating" (among peers) — pending structure. Where pickup of pending structure must be implemented by a return to a local root node, we must consider a rhetoric of re-arrival. In this case a component is not so much "departed" as "closed". (In this respect the Grønbæk-Trigg formulation for the distinction between composites and links cited above is right on the mark.)

Finally, it is interesting to relate the issue of pending structure to a protocol such as the Open Hypermedia Protocol (OHP). There is an interesting resonance between protocol design and rhetoric. Protocol designers often construct scenarios: idealized or paradigm exchanges between parties which serve as test cases for what messages need to be passed and what state rules should apply. There is an uncanny similarity between such scenario construction and the kind of exemplar studied in rhetoric. (For OHP scenarios, see [35].) Pending structure is analogous to protocol messages not yet sent. This may in fact provide an implementation framework for displaying pending structure. If an OHS component has a display mechanism for displaying the result of acting on an OHP message, perhaps a similar mechanism could be used for displaying a potential message which is "imminent" but not yet received, similar to a look-ahead mechanism that pre-fetches link targets not yet followed as a performance optimization.

GENERALIZED BOOLEAN MULTI-ACTEME RELATIONSHIPS

If actemes can in exist in "or" relationships and "and" relationships, then clearly we should investigate a more general logic of acteme relationships. The relationship of boolean implication is related to the concept of Guard Fields [6]. A guard field is a form of conditional linking in which a link is accessible only

after some other link has already been traversed. To be somewhat more precise, a guard field expresses a boolean relation

¬X É ¬Y

— if you haven't activated acteme X then you can't activate acteme Y.

The subject of oscillation has already been mentioned. Oscillation may be considered an example of boolean exclusive-or: object X is presented or object Y is presented, but not both.

Clearly we should be open to the use of the full range of boolean possibilities. A difficult question here is: what should show in the user interface concerning the boolean relationship of actemes? Typically guard fields are not displayed to the user. For instance, Afternoon [23] gives the reader no cue — visual or otherwise — that a guard field exists. It is not unknown in literary hypertext for critics to determine that behavior at a particular lexia is controlled by a guard field only by opening the hypertext in a full authoring version of the hypertext environment — an option not available to the typical reader. Where boolean behavior such as conditional linking is controlled by scripting — e.g. JavaScript in the case of the Connection System [25] — the script itself may be accessible to the reader; JavaScript is always accessible from the "View Source" menu command. Certainly viewing the source code for scripting is a less extreme act on the part of the reader than opening a hypertext in an authoring system, but even View Source may be described as a heavy-weight activity that should not be expected of the reader.

CONJUNCTIVITY IN THE LARGE

In what preceded, we have been considering what may be described as conjunctivity in the small: conjunctivity at the fine-scale granularity of actemes. In this section we consider much larger-scale hypertext activity — at least at the level of what was described in [39] as the episode.

CONJUNCTIVE NARRATION

The conjunctive, or additive, character of narration has been a subject of commentary beyond the realm of hypertext. Walter Ong, for instance, comments on the conjunctive character of biblical narration [34 p. 37]. A more complex case is the narrative method described by Genette as iterative [13 chap. 3]. Iterative narration collapses several repeated occurrences into a single passage. Genette gives extensive treatment to Proust's use of iterative narration. The iterative concept may be said to have an inherent conjunctivity, in that multiple occurrences are conjoined into a single description. Beyond that, because the reference of an iterative passage can span considerable amounts of time —— and therefore several "locations" in the plot of a narrative, there is a kind of implicit additivity to the overlay of multiple separate iterative passages, which act together in a kind of collage effect. Finally, for an interesting discussion of conjunctive cinematic narration, see [28].

Of course the sentence itself may be described as a conjunctive structure: the parts of a sentence, such the noun phrase, verb phrase, etc., are hardly alternatives to one another.

An interesting hypertext example of conjunctive narration is provided by Califia [9]. While in the small Califia's formal devices appear to be those typical of disjunctive hypertext, Coverley herself states that the concept of conjunctivity was explicitly in her mind as as she wrote this work [personal communication]. There are a number of specific effects at work in this hypertext that reinforce its conjunctive character. The use of photography to evoke specific times and places is very distinctive; each new photograph gives the reader the feeling of adding to an album of experience of the characters in the novel. The navigational interface continually presents paths for the characters that reinforce the importance that the story is the story of all of them. Indeed, throughout Califia the navigation is more additive than alternative.

There are several important questions about how the conjunctivity of narration is functioning in a particular hypertext: (1) Is it clear to the reader what pathways might be taken as additive as opposed to alternative? (2) For those pathways that the reader has taken as additive, does the addition "take"? (In the terminology from above, this is the question of whether conjunctions have been

successfully actualized.) (3) How have the reader's expectations about the amount of additivity been satisfied?

SECONDARILY CONJUNCTIVE HYPERTEXT

A hypertext may employ typically disjunctive relationships among actemes when viewed locally, but a more conjunctive picture may emerge from the reading experience as a whole. The most familiar concept illustrating this idea is the well known topic of contour [5]. In [24 pp. 82-83] Michael Joyce describes how a spatial view in Storyspace — and even more indirect structures such as "link-plots" — can give topographic overviews to hypertext. In such a view the density of pathways converging on a particular node can emphasize the additive character of those pathways, in a way that is not apparent looking at a particular lexia.

An interesting form of secondary structure which may be called the ***link-name lexia***[31] is found in ***Samplers***, by Deena Larsen. In this work a set of link names brought up by a menu forms a lexia in its own right; Larsen recites these lexia as small poems when reciting this work. Conjunctivity here is somewhat equivocal. These link-name lexia typically form sentences, and as noted above, the sentence is a conjunctive structure. However, it is not clear that the structuring provided by the link-name lexia is intended to structure the link-following actemes in their original source lexia.

The concept of gathering was discussed in [39]; in effect gathering is the construction of a secondary hypertext with materials carried from the primary hypertext being viewed. Even if the primary hypertext is purely disjunctive, when materials are assembled in a gathering interface, that interface may offer conjunctive methods, such as spatial hypertext. Structural methods from the secondary hypertext may be used to organize higher-level structures from the primary hypertext transparently. In this way the reader's hypertext may become conjunctive even when the author's hypertext is not.

Various forms of secondary texts are customarily employed in literary theory. Such concepts as plot and close reading are both examples of secondary

31 This term is mine, not Larsen's.

texts. Close reading poses a difficult question. A close reading may be described as a "discourse text" where there is an intense localized topological mapping between the discourse text and the text being discussed. It is at least arguable that by this measure, it is impossible to create a close reading of a hypertext without creating a second hypertext. However, we are severely lacking examples of this. Can one properly give a close reading of a disjunctive hypertext as a conjunctive hypertext, or vice versa? Or, should a close reading match the conjunctivity of the text being discussed? Plot is quite often discussed in additive terms; a scene is often said to add to (or to complicate, and in that sense subtract from) our understanding of the plot. Perhaps readings of hypertexts will devise a new body of practice with a new concept of reading not entirely matched by the close reading model. Answers to these questions await further study.

CONCLUSION

A great deal of hypertext may be summarized as endowing the word with activities that provide multiplicity in amidst the words from the reader's point of view. These multiplicities may provide for alternatives in the disjunctive case, or combinations in the conjunctive case. Conjunctive multiplicities raise many issues that don't occur in disjunctive hypertext, involving peer traverse among the terms of the multiplicity at the local level, actualizing the conjunction, and indicating what part of a structure is still pending, among others. "Building the transaction of the and" raises many fruitful issues for further study.

ACKNOWLEDGMENTS

I am deeply indebted to Randall Trigg for his extensive comments on a draft of this paper.

[7] The self-referential character of hypertexts with tabletops should be of great interest to the literary community, considering the widespread use of self-referentiality in postmodern literature.

[8] This term is mine, not Larsen's.

LOCUS LOOKS AT THE TURING PLAY: HYPERTEXTUALITY VS. FULL PROGRAMMABILITY

INTRODUCTION

It would be possible to implement a hypertext using purely physical means, without any recourse to software at all. Still, when one hears the word "hypertext," one thinks most often of text in which software plays a crucial role. Hypertext is a particular kind of software, however; hypertext is not coequal with all software. Hypertextuality is typically viewed as a property which derives from using such particular software. Hypertext systems have varied greatly in their degree of extensibility. Although hypertext systems have been fully programmable ever since their inception [16], [45], there is a wide perception that the advent of Java has brought about a change in atmosphere, in which generalized programs might be found in hypertext anywhere. (E.g. [41].) This paper attempts to address issues facing the identity of hypertext in the face of fully Turing complete generalized programmability. (By "Turing complete" is meant a system containing a programming language with sufficient power to emulate any theoretically achievable calculation; the name refers to an abstract machine devised by Alan Turing which he proved could emulate any calculation process.) Although the focus will be on literary hypertext, it is hoped the discussion will be applicable to a broad range of hypertexts. In discussing literary hypertext we will make a point of opening the discussion to cybertextual1 forms that may or may not be considered hypertext depending on one's point of view; to extend the discussion to more generalized kinds of algorithmic texts is important, since the issue is exactly the relationship of hypertext to more generalized algorithmic forms. A theme which will emerge throughout this paper is that dimensionalities pertaining to the algorithm are one way in which genre may be expected to operate in cybertext, and that hypertextuality—rather than being "all or nothing" is a condition to which only some of these dimensions may apply. Hypertextuality may vary (e.g. be artistically varied) even within a single work.

Since the subject is hypertextuality vs. generalized programmability, we begin with a brief survey of the kinds of strategies that have been used in hyper-

text systems to provide extensibility — i.e. to add generalized algorithms into an existing hypertext system.

A BRIEF REVIEW OF HYPERTEXT EXTENSIBILITY

Any hypertext system may be considered extensible if the source code is published. Of course publication of source code is rare in commercial systems. Apart from the ability of a user to adapt generalized algorithms to a hypertext system by modifying source code, the following strategies are among those that have been used to provide hypertext extensibility.

EXTERNAL EXECUTION

A hypertext system may allow calls to external routines, which may be in any programming language supported by the native operating system. Examples of this type of extensibility are the cgi-bin interface of HTML [23] and the HyperCard XCMD/XFCN interface [29]. This type of extensibility may be severely limited by the circumstances under which the call is allowed to occur. For instance a cgi-bin script can be executed from an HTML page only by clicking on a link; the cgi-bin mechanism cannot be used to extend HTML behavior to implement e.g. "no-click" hot-spots.

INTERNAL LANGUAGES / PUBLISHED INTERNAL PRIMITIVES

Some hypertext systems have been built on top of programming languages. NLS/Augment was built in a specially constructed language called L10, with interface components in a language called CML [46]. (And it should be noted that in NLS/Augment, links are specified fundamentally by addressing [17] which is basically a programmatic concept.) In this approach, extensibility derives from publishing to the user the primitives on top of which the hypertext system was built, allowing the user to add additional functionality using those primitives. NoteCards was built on Lisp primitives that the user could invoke [44]. Other hypertext systems, such as Trellis [19] have added a layer that includes an internal language offering a great deal of extensibility.

SCRIPTED ARTICULATION POINTS

A hypertext system may allow the author to work in units; at the points where these units "go together" the author may be able to interpose an algorithm in a scripting language supported by the hypertext system. Such languages are typically unique to the hypertext system, rather than being generalized operating-system-level languages. Examples include HyperCard [20], MacWeb [34], KMS [3], and JavaScript [24]. For instance HyperCard uses such units as stacks, cards, and buttons; depending on what the user does messages are generated; these units can contain scripts in a language called HyperTalk — a fully Turing complete language — which are triggered when the unit receives a message.

One weakness of this method of extensibility is that while the programming language may be fully Turing complete, it typically does not allow the author to change the structure of how the "articulation points" work. For instance in HyperCard, a HyperTalk script author may not alter the inheritance sequence by which objects receive messages. Or another HyperCard example: an object which receives a mouse-down event then "owns" the mouse, directly receiving all future mouse events until a mouse-up event occurs. This behavior cannot be altered through HyperTalk scripts.

GENERALIZED OBJECT INHERITANCE

Where a hypertext system is written in a fully object oriented programming language, such as Smalltalk, extensibility may be provided by the general mechanism of object inheritance provided by that language. Examples of such systems are Sepia [43], Dolphin [21], Hyperform[47], and HyperDisco [48]. (One could make the case that this is really the same concept as published internal primitives; the object inheritance is simply the method by which the primitives are published.)

GUEST ALGORITHMS

In HTML a Java applet maintains complete control over a window, where it becomes a kind of "guest algorithm". This can pose difficult aesthetic questions. There may be only a limited relationship between the guest algorithm and its host lexia; the guest algorithm may use a completely inconsistent interface from that of the host, and may "subvert" the host to the point of becoming a cuckoo that

takes over the whole nest. At this point the host framework becomes largely irrelevant, and we are left with a study of the guest algorithm as if there were no host.

Guest algorithms provide a way to experiment with "user interface laboratories" within otherwise inflexible systems, such as HTML.

IDENTITY OF THE USER INTERFACE / LOCALIZATION OF THE ALGORITHM

LOCALIZATION OF THE ALGORITHM

"Classical" hypertext might be described as text whose units bear with them their own user interface. Such fundamental concepts as following a link are typically triggered by operating a user interface object which in most cases visually appears to be in amidst the text as a visually marked anchor.2 "Widgets", such as a standard scroll-bar or dialog box, have locations where the user is to perform some action which are likewise clearly marked. The algorithms encountered when such user interface objects are operated may seem fairly trivial, but they may in turn trigger "handlers" for algorithms of great complexity (as in the scripted articulation points concept above.) While it is typical in the case of link anchors for these to have clearly articulated boundaries, the presence or absence of a handler, and the nature of this handler (if present) may have no visible indication in the user interface. A handler may impose conditional behavior, such as a Storyspace guard field [6]. In literary hypertexts containing guard fields, there is typically no visible indication at all concerning whether a guard field is present, or what its parameters are. (E.g. see Afternoon [26].) Hypertexts presented by means of HyperCard may allow no access at all to the source code underlying event handlers; thus while the user may encounter a button with clear boundaries, there is nothing clear about the nature of the algorithm triggered when the button is clicked. Thus while some might consider this poor design, it may not be possible to place a clear boundary on the concept of "user interface" and state clearly just what kind of algorithm may be involved.

One property often found in hypertext algorithms is that they are highly localized. Or, as Deleuze and Guattari [14] might have put it, the algorithm in hy-

pertext is typically territorialized. An anchor exists at a certain location in a lexia. User interface widgets are found at predictable locations on screen objects. Generalized algorithms, however, maintain state information which may be global to an entire system. Is localization part of what we intuitively consider a characteristic of hypertext objects when distinguishing hypertext from more generalized kinds of software? While a guard field is "more algorithmic" than a link with no guard field, the guard field concept is still highly localized: whether a link can be traversed is dependent on what other links have been visited. [28] implements what the author describes as "floating links" [27] in which a localized user interface button is connected to a global state variable used in algorithms to determine what text is presented to the user. In this case, user interface behavior somewhat resembling familiar links is used algorithmically in a highly non-local way which the user must simply "discover" to determine the effect of having performed this user interface operation. Similarly, the cybertext Book Unbound by John Cayley [12] uses a localized interface familiar from hypertext — clicking on a range of adjacent words — as input to its text generation algorithm, but the results are highly non-local: words clicked on one screen may appear in any future screen. (But one may argue that in this case there is a local character to the reappearance of the selected words when they do appear.) The HTML concept of "cookies" [24, Appendix D] was devised to overcome the stateless nature of HTTP transactions, achieving global state from a simple action of following a link.

Even in conventional hypertext, when the user follows a link, the resulting lexia may be generated by an algorithm. E.g. MacWeb [34] can produce lexia that are generated "on the fly". Such an algorithm is localized in the sense that the lexia is localized in the overall geography of the hypertext, though of course such an algorithm may have global state. A similar situation exists where a link is computed dynamically, e.g. a Microcosm generic link [18].

While hypertext may have had its origins in text with attached algorithms of a highly local character, shall we insist on this as a characteristic of hypertext? This seems arbitrary and unreasonable; one may say "global yearnings" are becoming increasingly common in working hypertexts. Localization is simply one among several dimensions of the algorithm which we will consider in this paper as emerging dimensions of cybertext genre.

THE USER/ALGORITHM RELATIONSHIP

A user interface object may be characterized as an algorithm with a clear and predictable behavior triggered by localized identifiable objects; the relationship between the user and this algorithm may be described as master/slave where the user is the master and the algorithm is the slave. By complete contrast, a purely algorithmic text, such as the poème animée "Soleil" by Patrick Burgaud [11], presents the user with almost no interface — indeed no interactivity at all; the user is simply present as the algorithm unfolds its results, much as the viewer is present at the cinema. (Though the user can quit.) This relationship between user and algorithm may also be described as master/slave, but in this case the user is the slave and the algorithm is the master. Unlike user interface algorithms, in this type of cybertext there may be no predictability at all to what the algorithm will do: one must simply discover this by observing its results — perhaps during several sessions (see the discussion of "sampling" below) — much as one must discover plot by simply observing a story unfold.

Other user/algorithm relationships are possible; e.g. in a game approach, the user might "play against" the algorithm, making the user and algorithm peers of a kind [13] (see also [38]).3 Of course in literary hypertext an author need not make a rigid commitment to a single approach; interface elements might sometimes be invisible, might sometimes be operated entirely under control of the algorithm, and this might change without visible indication to the user. For instance "passage" by Philippe Bootz [10] is an example of a literary cybertext in which the user has a certain (unknown!) window of time in which to act; if the user does not make a choice within that window the algorithm will act on its own.

The master/slave analysis of the user/algorithm relationship is similar to but not quite the same thing as Aarseth's discussion of determinacy [2]. A user interface device could trigger random behavior under "control" of the user but without predictable outcome (e.g. Judy Malloy's Its name was Penelope [31]), and conversely an algorithmic text could be completely determinate yet leave the user in the role of slave to the algorithm as master.

IDENTITY OF THE ALGORITHM

Where algorithms are confined to simple user interface elements, the identity of the algorithm is quite clear; it is established completely by (1) the boundaries of the objects that trigger the user interface behavior and (2) the specification of that behavior. It is typical for such behavior to be completely specified in documentation that explains how to use a hypertext. For other types of cybertext the identity of the algorithm is much more problematical. One possibility is that all algorithms are accessible to the user, and indeed are considered by the author as aesthetically integral to the work. Such accessibility might in fact operate by means of typical hypertext operations, e.g. offering the text of a handler algorithm as a special type of link [34], [20]. (Of course the algorithms of a cybertext might be completely inaccessible. Some degree of inaccessibility of the code is in fact typical; most computer usage takes place on systems for which accessibility of source code for the operating system or major applications is the exception. Source code for commercial systems such as Storyspace and HyperCard is simply not made available.)

Should the source code for literary cybertext be an aesthetic object? [13] gives a striking example of a cybertext where the algorithm is itself a poem. Such works are currently the exception.

SAMPLING ACTIVITY STRUCTURE

Where the algorithm is not accessible directly, its effects may be understood by sampling: observing the results of the algorithm in repeated sessions. (Aarseth [2] describes very briefly a somewhat similar concept as "playing for plot".) This poses some interesting issues for criticism. Consider the case of poems generated algorithmically, such as a work by Jean Pierre Balpe [4]. Formally, each poem has an appearance identical to a poem that might have been written "by hand"; because the algorithm is inaccessible, the only way to determine the aesthetic characteristics of the algorithm is by repeatedly sampling the poems. In this case we are somewhat removed from hypertext as it is usually construed: once the poems are generated there is no interactivity to them at all. But it would be a mistake to say there is no interactivity involved whatsoever in this work: it is the user who decides how many poems to generate, i.e. when to stop. Is the artistic work in this case (1) those poems actually generated (2) all poems which might be generated (3) the algorithm itself — even though this is completely hid-

den? Discovery of the algorithm through sampling is not so very different from discovery of the topography of a hypertext through the discovery of contours [5], [25]; the formal similarity of the poems produced by a generator algorithm such as those employed by Balpe to poems that might be written by hand is not so very different from the similarity of the individual lexia in many hypertexts to pieces of conventional linear text.

Where aesthetic issues pertaining to algorithm sampling may differ significantly from hypertext aesthetic issues of lies in the organization of the sampling: What is the structure of sampling activities? Hypertext activity is structured, as reflected in devices in the hypertext [40]. Repeated samplings of the results of a literature generator may offer no clear activity structure. On the other hand, where algorithms themselves are accessible by means of hypertext activities, e.g. links, algorithm sampling activity may be structured by the activity structure pertaining to these activities. The extent to which algorithm sampling activities are structured is yet another dimensionality of cybertext genre.

IDENTITY OF THE PROGRAMMER

Closely tied to the issue of the identity of the algorithm is the identity of the programmer. While the literature of hypertext rhetoric is replete with discussions involving the role of reader as writer (e.g. [30], [25]), far less attention is paid to the tripartite agency of reader/writer/programmer. (For discussion on this point see [39], [13], [38].) This is a difficult issue. Surely not all writers will relish the thought of becoming programmers. Should extensibility be extended to the reader? If we are to give the reader the freedom to participate in constructing a hypertext, it is arbitrary and unreasonable to impose an artificial boundary prohibiting the reader from participating in constructing algorithms in a more general sense. At what point does "authorial intention" reside in the algorithm?

STRUCTURE VS. BEHAVIOR

Hypertext analysis and rhetoric have long been concerned with structure; one may say the node link model is an inherently structural concept. On the other hand, it is characteristic of an algorithm that it exhibits behavior; underlying structure may be much more problematical. At its most extreme an algorithm

may exhibit "nothing but pure behavior" with no underlying structure at all. Thus the issue of what kind of algorithms we might call hypertext is deeply involved in the relationship between structure and behavior. In this section we explore these issues directly.

STRUCTURE VS. BEHAVIOR IN THE CLASSICAL NODE-LINK MODEL

The node-link model of hypertext — say as elucidated in the Dexter Reference Model [22] — presumes an underlying structure, namely the graph formed by the nodes and links. The system implementer is deeply involved in this structure, since the software comprising the hypertext system must maintain it and provide a way for the hypertext author to construct it. When a hypertext is complete, the extent to which this structure is accessible to the reader varies considerably with the particular hypertext. There may or may not be a graphical view attempting to give the reader a direct view of this structure, the hypertext author may encourage or discourage the reader from focusing on the structure, etc. Accessibility of the structure may vary among categories of reader; e.g. those readers with access to a full authoring environment for the hypertext system may have access to graphical views of the structure, while those with only a "run-time" viewer may not. (See [2] on this point.) Still, regardless of how directly the underlying structure can be accessed, the reader is aware that there is such a structure, and it heavily influences what a reader will do. For instance, if a graphical view of the structure is not available, one of the things a reader may attempt to do is form a "mental map" anyway [15].

Behavior in the node-link model — as experienced by the reader — is typically confined to navigation. When a link is followed, the system is expected to respond by presenting the lexia at the target end of the link. Other behaviors are related to choices of where to navigate; e.g. the user interface may be expected to bring up a presentation of what links are available, possibilities of backtracking etc.4 Behavior thus takes place explicitly with reference to the underlying structure.

Of course for the hypertext author, the authoring environment will offer many behaviors related to constructing the structure.

BEHAVIOR IN ALTERNATIVE HYPERTEXT MODELS

Various alternative hypertext models have been proposed, e.g. relations [32], piles [33], sets [36], Petri Nets [42], simultaneities [37]; some of these kinds of structure may be described as conjunctive rather than disjunctive [37] in that rather than viewing e.g. links as alternatives to one another, the user forms an abstraction consisting of the combination of elements. The behavior of the hypertext system is called upon to assist in this process. While it is not necessarily quite the same thing as navigation, such behavior is still highly focused on structure: the behavior is aimed at bringing the reader to "construction points" in the structure.

THE STRUCTURAL POINT OF VIEW

Nürnberg, Leggett, and Schneider [35] presented a view of hypertext as just one example of what they call structural computing. In this paradigm, hypertext concepts are reformulated in terms of generalized "structure stores" and "structure processors". Behaviors are abstracted separately, and viewed as "computations over structure". This paradigm makes explicit the primacy of structure, which it seeks to generalize broadly to many realms of computing beyond hypertext. When working from this point of view, the question of "what behaviors are hypertext" seems strangely irrelevant. The degree to which a system should be considered hypertext would logically focus on the nature of the structure stores and structure processors; presumably any behaviors of such a system would inherit hypertextual characteristics from the nature of the structures they operate on. Under this paradigm, there is an abstraction layer for behaviors, but behaviors are to operate on an existing layer of structure stores.

How would a generalized algorithm fit into this scheme? While virtually any of the extensibility strategies discussed above would "work", generalized behavior not in accord with the structural framework might pose difficulties. Probably the scripted articulation points or guest algorithm concepts would be the easiest to implement.

THE BEHAVIORAL POINT OF VIEW

The opposite point of view is also tenable: a primary focus on behavior, with no preconceptions about structure. In a number of papers poet Philippe Bootz has argued forcefully for the functional point of view [7], [8], [9].5 Figure 1, reproduced from [9], shows part of his scheme. The textes-auteur consists of notations prepared by the author for the programmer who will implement the génération — the "algorithm box"; these notations might be such materials as paper scripts or storyboards, i.e. not necessarily machine readable. The texte-à-voir is the textual layer accessible to the reader. The texte-à-voir appears based on whatever functional devices trigger in the "algorithm box" (génération); structure within the algorithm box is not generally accessible. (The layer shown as texte-lu — "text read" is a mental construction layer created by the reader; this is somewhat analogous to the notion of gathering as presented in [40].) One should note that Bootz's poetics differs considerably from much of the rhetoric familiar in the hypertext community: contrast his insistence that the various domains of author, text, and reader be separate vs. frequent assertions of reader/writer interchangeability [25], [30]. (There is no feedback loop in Bootz's scheme from the texte-lu to the textes-auteur.) Or, consider his concept of l'oevre verrouillée ("locked work") vs. constructive hypertext [25]. In this framework, structure is entirely contingent on what happens in the texte-à-voir.

What type of extensibility might open a hypertext to the behavioral point of view? Obviously the guest algorithm concept comes to mind, but this has the difficulty discussed above of cognitive dissonance between the native hypertext behavior and that within the guest. A more natural approach is the published internal primitives concept. This would allow the behavioral point of view to "be in charge" and yet use hypertext behavior where appropriate by simply invoking it.

DIMENSIONS OF HYPERTEXTUALITY

The common view of hypertext is that one chooses to work in a particular hypertext system, and the result becomes hypertext — "thereby". The picture for cybertexts that allow full generalized programmability is much more complicated. In such a context, there is no reason whatever to assume that hypertextuality is "all or nothing". Rather: there are dimensions to hypertextuality; these dimensions become artistic variables, just like other artistic variables. Some could be present with others absent; the author might vary completely the degree to which

a dimension is present depending on where one is in the work. (For instance, in "passage" [10] a mouse cursor — in the image of a computer mouse — can appear at some points; when this cursor appears the user can click and obtain hypertextual behavior. One doesn't really know how or where or when this cursor might be available; sometimes it happens, sometimes it doesn't.) How an author treats dimensions of hypertextuality is one of the ways that genre may be expected to emerge in cybertext space. Let us now review these dimensionalities:

- Localization of the Algorithm

Typically hypertexts have been highly localized. The text occurs in units, i.e. the lexia, of material presented at one time; locations where algorithms activate, such as link anchors, are clearly marked. Algorithm behavior takes place with respect to these localizations: events are triggered by localized activity; these events typically change one's location in the hypertext. A generalized algorithm can contain as much or as little localization as the author wants; localization can vary depending on past user input and the current state of the algorithm.

- Degree of Algorithm Identity

Classical hypertext algorithms have a clear identity: the user knows what is supposed to happen; indeed it would be taken as a sign of bad design if the user were not to know what is supposed to happen. But in the literary world, incomplete knowledge on the part of the reader has been an age-old artistic variable — the novel derives much of its power precisely from the fact that the reader doesn't know what is going to happen. In generalized cybertexts it may be artistically important for the author not to spell out the identity of the algorithm. The author may or may not want the algorithm itself (e.g. source code) to be accessible; the author may or may not want the reader to know whether a particular phenomenon occurred as the result of an algorithm.

- Structural Focus vs. Behavioral Focus

Hypertext has typically been a domain with a high degree of structural focus. In the node-link model, the graph giving the link relationships is a structural

graph; great attention has been paid to issues of how to convey this structure to the reader. A generalized algorithm exhibits behavior; this behavior may or may not clearly resolve to a structural background.

- User/Algorithm Relationship

The relationship between the user and the algorithm in hypertext is typically master/slave with the user being the master and the algorithm being the slave; in generalized algorithmic cybertext, any user/algorithm relationship is possible. It could be peer-to-peer or master/slave where the user is the slave, or a complex combination. As above, this can vary within a single work depending on the state of the algorithm.

- Activity Structure of Algorithm Sampling

Where the algorithm itself is not accessible, the nature of an algorithm may only be revealed by exercising it repeatedly. These different sampling events may or may not have an activity structure; if they do, it may or may not relate to dimensionalities explored above. E.g. it may or may not have "topographical" identity with respect to localizations. The author may or may not give guidance on how to do sampling. If there is an activity structure it may resemble hypertext activity structure even if there is little resemblance to hypertextuality along other dimensions.

ARCHITECTURAL SUPPORT FOR ALGORITHM GENRE

A system designed to provide support for hypertextuality yet be open to the full range of possible cybertext algorithm genres faces some interesting challenges. The guest algorithm concept would certainly support any possible algorithm, and hence e.g. any possible approach to behavior vs. structure. For the cybertext author, however, the guest algorithm concept is extremely stark: it presents the author with a "blank page" programming concept — i.e. no real architectural support at all. The scripted articulation point concept does provide for somewhat flexible behavior, though within the confines of the architecture of the articulation points. (It should be noted that HyperCard is far more popular as

a cybertext authoring tool among poets — who often need more flexible behavior — than among fiction writers.)

From the point of view of a cybertext author, the most desirable approach to extensibility would be to blend all of the extensibility strategies mentioned above and make them all available. The scripted articulation points concept can be achieved on top of published internal primitives. The concept of a guest algorithm space can be offered to guest primitives — ideally presented as some form of construction kit. A construction kit built on top of published internal primitives would offer off-the-shelf abstractions to those who need only a modest amount of extensibility, yet provide all the flexibility needed by those with more extensive programming requirements.

ACKNOWLEDGMENTS

I am grateful for the comments of Catherine C. Marshall, who read an early draft of this paper.

1The term 'cybertext' will be used more or less in accordance with [1] as a generalized form of machine mediated text without being specific about what type of machine mediated text. Cybertext comprises hypertext, non-interactive machine generated texts, MUDs, etc.

2I am indebted to Rosemary Simpson [personal communication] for pointing out that the visual appearance of embedding of an anchor in the lexia is completely independent of whether the hypertext system embeds links within the text or maintains them externally; this implementation issue is not addressed in this paper.

3The peer-to-peer relationship between user and algorithm is even more pronounced in MUDs and MOOs. We will not discuss MUDs in this paper. For an extensive analysis of MUDs as cybertexts see [1].

4Behavior with respect to composites is still an unresolved issue.

5Regrettably, the hypertext community has not been aware of this work; most of these papers originally appeared in journals not likely to come to the attention of hypertext researchers. There is almost an entire parallel literature on cybertext in French that has intersected poorly, by and large, with the hypertext literature; for a review of the French work see [7].

A HYPERTEXTUALITY OF ARBITRARY STRUCTURE: A WRITER'S POINT OF VIEW

INTRODUCTION

Language is intensely structured. At a fine grain it is structured by syntax; at a still finer grain by phonology; at a coarse grain by such structures as argumentation and rhetoric; at a still coarser grain by plot, narration, etc. Hypertext brings something new to the table: explicit external structure, coupled with interactivity wherein that structure plays itself out. This paper takes as a given the premise that an artistically open hypertextuality must allow its concept of external structure to be as arbitrary as possible, notwithstanding the extent to which highly particular models of structure have become extremely popular. A number of issues will be analyzed from the point of view of openness to arbitrary structure; this analysis will be strictly from a writer's point of view rather than from a system implementor's point of view, though of course it is hoped this analysis will be useful to system implementors. In particular we will take no position at all concerning such issues as how structure is managed by a hypertext system (e.g. as an explicit "structure store" layer [16].) We begin with an analysis of what "structure" entails. (To some the idea of what constitutes structure may seem obvious, but it bears investigation.) We will be at great pains to avoid constructing Yet Another Formalism; not because any doubt is being cast on the worthiness of that activity, but because we fear the act of constructing a formalism risks prejudging the outcome of what will constitute structure. Although structure is the subject at hand, a great deal of emphasis in this paper will be placed on behavior, for a number of reasons. It is typically in the area of inflexible behaviors that off-the-shelf hypertext systems most disappoint artistically inclined authors. Moreover it is all too easy to allow behavior to slip still further from view in the rush to place emphasis on structure. From an artistic point of view, the key word is openness: as we explore how to make hypertext open to the widest variety of possible structure, it must also remain open to the widest variety of possible behavior — and the widest variety of relationship between the two. Authors will simply not put up with hypertext systems in which behavior is the poor handmaiden of structure.

EXTERNAL STRUCTURE

CONCEPTUAL OVERVIEW

In saying that hypertext consists of words (or other media) with an external structure, "external" is meant with respect to the words themselves. While it would be possible to indicate hypertextual structure entirely with words (including numbers) contained in a text — and in fact exactly such analysis is done to situate historical antecedents of hypertext (e.g. [3], see also [14]), typically a hypertext uses means external to the text to contain the structure. This concept of external is not to be confused with an external concept of storage, e.g. an externally maintained linkbase vs. embedded links in the data sense. The most widely discussed structural model is the node-link model, e.g. as elucidated in the Dexter Hypertext Reference Model [7]. Many other modes of structure have been discussed including sets [17], relations [12], Petri Nets [22], simultaneities [18], and piles [13]. Articulation of how elements of the text participate in such a structure is part of The Rendering Problem, which is discussed below. With respect to such structural elements — as rendered — the user performs activities, which result in Behavior — also discussed below. In this section we discuss aspects of structure itself.

ASPECTS OF STRUCTURE

Before proceeding to aspects of structure that may be more familiar to those from a systems point of view, it is important to emphasize a particular aspect of structure often important artistically but neglected by many hypertext systems: the Null Structure. In a null structure, elements are presented with no structural relation whatever except that they happen to be juxtaposed. Elements in a null structure are simply there; it may not be pertinent to raise further issues concerning their "thereness". Similar to the concept of the null structure is Emergent Structure [13]: an informal structure perhaps constructed on a temporary basis which may serve as the basis of a more formal structure to be constructed later. VIKI spatial collections [13] are an example of such a structure.

Let us now look at some specific aspects of structure. While we eschew formalism, the discussion that follows is a kind of informal metalanguage for

the type of framework in which a structural description would need to be constructed.

SCOPE

A structural element may have a specific scope, indicating what other elements in the structure it relates to. An example is a set where the scope indicates which elements are members of the set; another example is a bounded region in a spatial hypertext where the scope is whatever elements fall within the region. Hypertext anchors are often conceived of as such a bounded region.(A set might be indicated as a region, of course, or set membership might need to be specifically designated independent of any presentation issues.) Yet another example is the familiar Dexter concept of composite, whose scope is all the elements contained in the composite. All XML tags are either designated as "point markers" or they have a clearly designated scope.

Some structural elements may be deemed not to have an effective scope. What is the scope of a hypertext link? There is room for considerable divergence of opinion on this. Some may consider the scope of a link to be the lexia or anchors at the two endpoints. Others may conceive of the link as being more like a programming GOTO statement and thus scopeless [4]. Any concept of structure must designate some elements as atomic; these elements may be considered scopeless (or such an element might be considered as having a scope consisting only of itself.)

VALENCE

The term valence is here borrowed from chemistry. Valence is a schema which represents for a given kind of structural element what must be "plugged into it" to have a valid instance. In many systems the valence of a hypertext link may be described as an ordered pair consisting of a source anchor and a destination anchor; in other systems the link may be much more complicated including perhaps a type (e.g. [15]), a link name, and so on. There may be some overlap between the concepts of scope and valence.

Depending on the structural framework, a valence specification may or may not include type specificity. E.g. a structural framework may require that only anchors may be attached to a link; a different framework might allow any structural object to be attached to a link. There might be severe restrictions on the type of object that is allowed to be plugged into a link name slot (e.g. practically no hypertext systems allow the name of a link to be a link.)

DIRECTIONALITY

For some structural elements, such as a set, the concept of directionality does not apply; for others it may be extremely important, e.g. the typical unidirectional hypertext link. In other cases directionality may not be formally specified but may be "metaphorically present" — e.g. an Aquanet relation may be deemed to relate its elements "together", implying an inward directionality toward some center. A structure such as a pile where the elements are so closely spaced that they overlap one another may have an implied outward directionality of all elements to spread themselves out.

BEHAVIOR SPECIFICITY

Behavior is a large topic, which will be discussed in more depth below. It is included here since a structural description may include behavior differences as the only way that structural elements are differentiated. E.g. in some literary hypertexts, "hot words" may not be graphically differentiated from those which do not serve as link anchors; the only way to tell is to click on the word and see what happens. (E.g. [9]).

The "dimensions" just given provide a kind of micro-language for describing structure; these aspects can be combined in arbitrary ways to describe a vast variety of structural possibilities. Openness to novel ways of combining such aspects is exactly what a hypertextuality of arbitrary structure is about.

LOCALITY/ATOMICITY OF STRUCTURE

Hypertext algorithms are typically localized [20]. Localization is also an issue that is important in discussions of structure. We spend much of our time in hypertext talking about highly localized structures, such as links. In fact, these

structures are often referred to as if they were atomic, when that may or may not actually be the case [19]. What level of structural locality is really relevant? An example may help to clarify this issue. It is sometimes asserted that Aquanet relations can be expressed in a Dexter formalism as composites [Randy Trigg, personal communication]. Presumably this would work by defining an Aquanet relation as a composite of links. Putting aside the difficult question of whether the behavior would be correct in this "translation" of the relation concept (this question is discussed below), it is important to note that such a composite must function as a complete unit; the individual links in such a composite have no particular meaning with respect to the original abstraction, namely the relation. The "locality" that is significant is the entire composite. Likewise, consider a structure where a link name is associated with its link by a special kind of link. Does this link have any independent existence of its own?

While such questions might seem to be splitting hairs, the same type of issues are considerably more thorny when projected up to higher structural levels. Consider the structural concept of neighborhood. Let's suppose that neighborhood is defined in graph-theoretical terms based on a node-link framework. E.g. the neighborhood of a node is defined as the set of nodes reachable via one link inbound to a given node. This may be very easy to define, yet difficult or impossible in practice to compute. Consider, for instance, the difficulty of computing neighborhoods where the hypertext system in question is the World Wide Web. In this case there is simply no hope of computing neighborhoods in any reasonable time.

While hypertext authors must be reasonable in the face of computational difficulty, it can be very frustrating for authors to be forced to work in a hypertext system where the system designer's concept of structure is rigidly local, the author has requirements involving higher-level structural units, and the ability to manipulate such units should be doable but isn't. Locality is a matter of context.

With this as a backdrop, we can consider some specific aspects of locality particularly as they relate to higher-level structure.

STRUCTURES MAY BE COMPOSABLE

Higher level units may be constructed by iterating lower-level units. E.g. the concept of neighborhood above is constructed as iterations of the link. Networks may be composed by joining them with links. Such acts of composition, while they may in the general case produce an unacceptable computational load, may be in fact acceptable in the particular context of interest to a specific author or user. It is a difficult question whether authors need to be "protected from themselves" to be spared such computational minefields. In this author's opinion, system designers must be prepared to include in their systems structural concepts of unknown computational load; determining whether a structural concept has "hit the limit" of what is computationally reasonable is part of the artistic process in the cyber-arts. In the case where a higher-level structural unit is computationally defined, e.g. by composition, of course the user must have the ability to interrupt the computation.

This results in an interesting question. What should be done with the partially computed structure when a computation is interrupted? There are two obvious possibilities: (1) the interruption should be considered an abort (perhaps a rollback in database transaction terminology) and the resulting structure delivered to the user (if at all) empty. (2) The structure should be delivered to the user in whatever degree of completion was achieved before the interrupt. Both choices pose difficult issues. One can imagine a concept like the neighborhood above where the user places an informal time constraint on the computation time, e.g. "let me see how much of a neighborhood I can construct for this node on the Web by letting this thing run overnight." In this instance the user is not served by having the interrupt produce nothing, even when (say by consulting numerous search engines) the system did in fact produce a significant number of inbound links. On the other hand, if we accept alternative 2, we are forcing the system designer to include in the system provision for incomplete structures. This could cause significant difficulties if a structural concept has consistency constraints that are not met.

STRUCTURE LOCALITY MUST BE SEMANTICALLY RELEVANT

Regardless of how the primitive structural concepts of one hypertext system can in principle be represented in the primitives of another hypertext system, this representation must have real meaning. Consider the case dealt with above

where Aquanet relations are represented as a composite of links. Does this unit function as a whole in the same way as the Aquanet relation? Does it have some identity distinct from other kinds of composites? Surely this would be required if the representation of relations as composites is more than a technical formalism: the relation concept has a semantic identity in the original hypertext system, and this identity must be preserved in the target hypertext system.

Or, consider an "index" constructed as a simple list of outbound anchors — say in HTML. How does one designate that this list functions as a unit? Such problems are quite easy to solve in XML: one simply invents (or better yet deploys some already agreed-upon convention for) a tag for this unit, and the scope of the tag designates the semantic locality of the structural concept.

STRUCTURE LOCALITY MAY NEED TO BE DESIGNATED EXPLICITLY

Hypertext systems do not always provide the proper flexibility to do this. Consider a node-link hypertext that allows bi-directional links. What is the scope of a reverse link? Can we assume in all cases that the scope of the reverse link is identical to the anchor of the forward link? Perhaps not. Consider a simple citation link. The anchor for the forward link might be nothing more than the citation notation in the style sheet of the document, e.g. a string that might look like "[47]". But depending on the circumstance, the scope of the reverse link might be an entire sentence, an entire paragraph, or possibly an entire section of a document. To make that scope identical to the anchor of the forward link might be extremely confusing for readers. In general, there may be no formal way to infer the scope of this reverse link; the author must simply designate it.

STRUCTURE LOCALITY MAY NEED TO BE DISCOVERED

This point is really the opposite of the point above. Discovery of structure by the reader may be an important part of an author's concept; explicit designation of structure may in fact defeat this purpose, since an author may intend that a multiplicity of possible "derived structures" should be inferable from the author's material. Structure locality might be defined not by the document author, but by a kind of annotative process by the reader. The hypertext literature is full of discussion of the reader acting as writer (e.g. [10]). Most often this is phrased in

terms of the reader authoring hypertext structures, such as links or annotations. But there is no reason to assume the reader might not want to amend an author's already-existing structures; for instance I might decide I as reader don't like the scope of an author's anchor and want to change it.

This raises an odd structural paradox. How is an author to indicate incomplete structure? E.g.: an author wishes to indicate that the scope of a reverse link is to be left to the reader, but within certain boundaries. How should this be notated? What are the interface requirements? Or, a similar example: the indeterminate link. Consider a case where an author wishes to indicate a link, where the target of the link will be determined in the future by some activity of the reader. (This could be based on a computation of which among many alternative lexia the reader visited for the longest time, some explicit action by the reader — say the first among alternative lexia to be bookmarked by the reader — or any number of other possibilities.) One method of handling such concepts of incomplete structure is to not indicate an incomplete structure at all until it is completed. This certainly simplifies the problem: we are now simply faced with a dynamically constructed hypertext. But surely this will not suffice for all authors. Returning to see whether the structure has been completed, and if so what it looks like may be part of an author's intention; in this case there must be an explicit indication of the incomplete structure.

There is an obvious advantage in indicating incomplete structure to those systems where structure is achieved by some form of formal language, e.g. a markup language; in this case an incomplete structure can be indicated by means of e.g. markup where parameters are missing, or by variables which have not yet received values. This risks, of course, that a parser might return error messages for this incomplete markup. To avoid this problem, incomplete markup would have to be encapsulated in some way so that a parser would ignore it until it was completed. For instance, completions for structural properties might be indicated by a special kind of link; when the markup was completed the link would disappear and the completed structure would replace it.

THE RENDERING PROBLEM

By "the rendering problem" is not meant simply the issue of rendering e.g. text and graphics in the usual sense of say browser rendering engines. Here we are concerned with explicitly structural issues. In the following section we will give a taxonomy of some of those issues.

PARTICIPATION IN STRUCTURE

Every user of the Web is familiar with distinct rendering for link anchors (outbound, of course!) Typically an outbound link anchor is underlined. Rendering of structure participation indications for a hypertextuality of arbitrary structure is much more complex. The issues here are related to ones discussed above, e.g.:

SCOPES, VALENCES, BEHAVIOR EXPECTATIONS, ETC. NEED TO BE EXPLICITLY INDICATED

Aquanet allowed valence to be rendered via explicit diagram concepts; alas this extremely valuable insight has not been carried forward into many current hypertext systems. The idea currently popular in HTML that structure participation can be rendered solely by means of "in-line" text attributes seems particularly shaky. To accommodate arbitrary structure, it is hard to see how we can avoid rendering by some overtly external means.

Rendering of scope is a nasty problem. There are many intuitively straightforward ways of rendering scope — provided we all agree that scopes are nested and don't overlap. Unfortunately that assumption is simply untenable if we are serious about the word 'arbitrary' in speaking of arbitrary structurality; the assumption that all scopes will be nested is just as outrageous as the assumption that all structures will be linear.

One approach to this problem is to have a generic method of rendering structure participation which is not explicit about the structural details. Consider the typical HTML rendering of an outbound link anchor as typographically distinct (in this case by underlining). A hypertext system could have a special typographical rendering that simply means: these words are structural participants. Obtaining the details might be achieved via behavior. For instance, an ex-

plicit structure indicator — perhaps an Aquanet-style diagrammatic rendering — might pop up as the cursor is moved over the structure participant, somewhat in the fashion of Zellweger's fluid links [23]. Scope could be indicated by a simple marker that denotes a scope boundary; the actual extent of the scope could be indicated by behavior.

RENDERING BEHAVIOR HAS CONSEQUENCES — AESTHETIC, HUMAN FACTORS, ETC.

This is a very difficult issue. Regardless of the advances in computing technology of recent decades, certain aspects of the computing environment are still scarce resources, e.g. screen real estate. It is inherent in the nature of hypertext that it allows structures of vastly more complexity than are easily represented by the geometry of the words themselves; the rendering of this structure has to fit somewhere. There are many conflicting demands upon screen real estate. Among them are:

- Maximizing the amount of "lexia text" which can be displayed readably
- Preserving on-screen informal or implicit structure
- Human-factors concerns regarding intuitive information display
- Aesthetic concerns per se of the author

How does a hypertext system allow rendering of structure participation without foreclosing an author's options in the face of these competing demands?

There are similar issues concerning time as well as space. One method of solving rendering problems is to use animation. (E.g. [23].) However, animation has time consequences: it structures time in a particular way. Depending on the multimedia contents of a hypertext, the time requirements of animations used to render structure participation could conflict with time requirements the author has regarding other issues.

One approach to solving these problems is "pluggable" behaviors: the hypertext system could provide a "construction kit" whereby an author could build rendering behavior from a palette of tools sufficiently rich that conflicts like the ones discussed above could be avoided. This is apparently the suggestion made in

the Structural Computing framework of [16]; however that framework leaves a bit vague the relationship between structures and behaviors.

In fact, it has been typical of most hypertext systems that they "hard-wire" many forms of behavior, specifically including the rendering of participation in structure. Hard-wiring behavior is just as restrictive as hard-wiring structure itself; a hypertextuality which is open to arbitrary structure must be open behaviorally as well. Of course this raises the issue of "arbitrary programmability"; for a full discussion of this issue see [20].

REPRESENTATION OF STRUCTURE

It goes without saying that a hypertextuality of arbitrary structure must have some way of rendering that structure. Hypertext lends itself quite naturally to very complex structures; it stands to reason that hypertext authors will want the best available techniques for representing complex structures. Ideally, as new techniques evolve, there should be some method of retrofitting them onto existing hypertext systems. [21] describes how multi-focus fish-eye views were retrofitted onto VIKI. That experience is, alas, not entirely encouraging for hypertext authors, in that it required the system implementors to do the work that would allow including fish-eye views. Ideally such work should be doable by document authors. To allow for this, hypertext systems will have to provide much more open APIs than is customary at present.

Representation of structure has most of the same issues discussed above under participation in structure; there are a few other new issues at this level.

SUBSUMING (COLLAPSE/EXPAND) STRUCTURAL ELEMENTS

Participants in a structure might be structures themselves; a rendering should allow the reader to subsume such a structure to a single entity and collapse it. Note that this is entirely separate from the kind of rendering behavior that might occur as a result of navigation. (Navigation is discussed below.) Consider a node-link hypertext which is tree structured in the vicinity of a particular node; let's assume the user is "located" at that node. Without moving from the node, the user might wish to collapse all "child nodes" to a single graphic entity each,

regardless of how many "grandchild" nodes (and further generations down) extend from a given child.

More generically:

RENDERING OF STRUCTURES MAY NEED TO INCLUDE FILTERS

The fish-eye view concept [5] may be thought of as a kind of visual filter; others have been proposed. Such filters might have options that govern how distortions are handled. (See [21] for a discussion of of how proximity distortion vs. alignment distortion was handled in fitting fish-eye views to VIKI.)

REPRESENTATION OF LOCATION WITHIN STRUCTURE

This is, of course, the celebrated navigation problem. A full treatment of this problem is beyond the scope of this paper. Suffice it to say that the level of generality required to support arbitrary structure will require an interdependent methodology for representing location in structure. Beyond the usual issues of navigation, a hypertextuality of arbitrary structure has some special issues pertaining to location:

LOCATION NEED NOT BE A UNITARY CONCEPT

Most hypertext systems have some kind of concept of "current location"; typically this concept is unitary. I.e. there is just one "current location". Consider, however, a structure such as a "symmetric" ternary relation. I.e. the structure is a triple, where all three slots are equivalent. One concept of location in this structure would assume that the reader will read the items occupying the three slots in some order, each one in turn; current location is a straightforward concept in this case, designating which of the three slots the user is reading. However, perhaps the author's concept is that the three elements occupying the relation slots should "play simultaneously". The concept of current location is now much more complex. Perhaps now all three slots are simultaneously open, and the user is located in all of them: a non-unitary concept of current location.

Many familiar user interfaces open a multi-paned window, or perhaps several windows simultaneously. Today's user interface may have a unitary bias

driven by the fact that there is only a single mouse cursor, but this unitary bias is certainly not logically necessary.

NON-UNITARY LOCATION MEANS RENDERED OBJECTS HAVE DEPENDENCIES

If the user is "located in" multiple simultaneously open rendered objects (e.g. Windows), it is quite likely that those objects will contain informational dependencies. Thus the hypertext system will have to provide the document author with a method for managing these dependencies. Unfortunately, this is a complex subject. Prior experience with frameworks for dependencies among multiple open objects includes such celebrated cases as the well-known Model View Controller paradigm (MVC) [11]. Experience with MVC is best described as mixed; Smalltalk environments with off-the-shelf classes for MVC have proved a rich ground for experimentation, but MVC has not been popular in commercial systems, to the point that it is probably not an exaggeration to describe MVC as obscure among today's software developers (never mind users!) Unfortunately, it has not been replaced by a comparable framework. To place the burden of inventing such a framework on today's system implementors is asking a lot; but it is asking even more to place this burden upon the document author!

PARASTRUCTURE

As discussed above, arbitrary structure must include informal and emergent structure; among the devices used to represent such structures are proximity, alignment, typographical conventions, and so on. As discussed above, parastructure may act as a constraint on other aspects of a hypertext system, in that structure rendering may need to preserve the parastructure.

THE BEHAVIOR PROBLEM

Behavioral issues have already figured prominently in this discussion; we now discuss behavior directly. The relationship between structure and behavior as it affects hypertext programmability has been discussed in [20]. This discussion will adopt that framework, with a particular emphasis on openness to arbitrary structure.

"PURE BEHAVIOR"

Analogous to the null structure discussed above is behavior with no structure at all. (See [20] for references to examples.) Pure behavior is simply "there"; it unfolds like a cinema and has no internal structure that the user can access. An example is an internally undifferentiated multimedia object which simply plays; perhaps it functions as a plug-in, or a guest algorithm in the terminology of [20]. Even though such an object may exist as a node in a larger structural framework, from a structure point of view it exists as an atomic object. Nevertheless, there may be some subtle interactions between such an object and other components that are specifically tied to points in a structure.

A PURE BEHAVIOR PLAYER MUST BE A COMPATIBLE GUEST

As noted above, arbitrary structure may require a dependency mechanism among "parallel" pieces. A pure behavior piece must be compatible with such a scheme. If it is controlled by a multimedia player object, perhaps that object is sent parameters or signals from other components in the structure. There must be a framework for how this takes place.

PURE BEHAVIOR MUST BE PROPERLY EXPORTABLE

One of the obligations of a hypertext system is to properly save its state at points where this makes sense — both for reader and writer. A pure behavior object must participate in this. Of course it might be an aesthetic decision on the part of the author that a pure behavior object should replay from the beginning whenever its state is restored, but another author might include behavior objects which take so long to play out that (intentionally) the entirety of the play sequence cannot possibly be completed in a single session; in this case it may be aesthetically important to the author for the object to resume where it left off.

Similarly, a pure behavior object must be able to participate in transitions. Transition behavior will be discussed below; in that context pure behavior is simply a special case.

"STRUCTURE BEHAVIOR" MUST ALLOW FOR PURE BEHAVIOR

Many hypertext systems assume that their own "built-in" behavior is the only kind of behavior which will occur. This can pose conflicts if a pure behavior object is inserted into the space. Consider, for instance, a user interface — perhaps one to assist the blind — which relies on sound. A player object has a sound track. What should be the relation between these? Should they play simultaneously? Should the user interface sound interrupt the "pure behavior sound"? Perhaps the user interface sound might behave differently if it knew that a "pure behavior" player object with sound was playing.

NAVIGATION BEHAVIOR

Navigation behavior is arguably among the most important kinds of behavior in hypertext. The expectation when a user e.g. clicks on a link is that something will happen. It is typical for hypertext systems to make assumptions about navigation behavior — assumptions with which the document author may not agree. An important issue in a hypertextuality of arbitrary structure is how the document author may create her own effects in navigation behavior. In the literary world we care a very great deal about these matters.

TRANSITION EFFECTS

Consider the simple act of following a link in HTML. It is customary for the browser to repaint the screen completely to render the target page. This can cause a noticeable blink if the replacement page is substantially the same — even at a high bandwidth. Environments such as HyperCard [6] allow for explicit transition effects that delineate how the screen will change from one "card" to another. Perhaps the user wants a smooth fade from one screen to another. Note this interacts with the issue of exportable behavior just discussed. If there is to be a smooth transition from a static "structure object" to a dynamic "pure behavior" object, enough of that pure behavior must be exportable to the system to accomplish this. This is a difficult problem if not explicitly designed in to a framework; it is hard to imagine how an arbitrary guest algorithm exports enough of this behavior to allow for transition effects.

Other kinds of transition effects include flying (e.g. Web Squirrel [2]) — navigation in Web Squirrel causes the reader to "fly" over the space to the target

object — and zooming (e.g. Pad++ [1]). It took considerable effort to implement the zooming behavior in Pad++. What is the poor author to do if both the kind of zooming behavior implemented in Pad++ and structural concepts not implemented in Pad++ are desired?

INTERFACE BEHAVIOR

Interface behavior has an intimate relationship to structurality. Consider a structure such as a ternary relation. Such a structure is known in advance to have 3 slots; an authoring system might use this information behaviorally in the case where, for instance, not all of the slots have been completed. E.g. when composing the relation, if the author attempts to "close" the relation in an authoring mode and not all slots have been completed, perhaps a warning dialog should be displayed. A reader would expect such behaviors as: open all slots simultaneously, close all slots simultaneously, display a list of the slots, etc. Some of these behaviors may raise difficult implementation issues in which the author may want to participate. Consider the case of opening all slots of a relation simultaneously. Screen real estate is always in short supply. Something will have to be hidden to open new objects. What should the "hiding policy" be? How can a document author participate in this policy? How can the reader? Let's suppose further that our ternary relation brings together elements from "distant" places in the hypertext. What should the behavior be when these places are opened? It is exactly issues such as this that raise serious difficulties for the kind of formal "emulation" exemplified by the belief that e.g. Aquanet relations can be modeled as composites of links. If a relation is simply a composite of links, the assumption must be that "opening a relation slot" consists of following one of these links — with whatever behavior is entailed by following a link in the system in question. The idea of opening all the slots simultaneously — in some form of coherent display — is not available. Thus the interface behavior of a relation may not be correctly modeled by the interface behavior of a composite of links.

PARASTRUCTURAL BEHAVIOR

Parastructure was discussed above; while parastructure is often associated with issues of spatiality, there are behavioral issues here as well. A simple hypertext link may be thought of as (in some cases at least) at attraction between the

linked elements; the opposite concept is also possible, namely a repulsion between elements. An author might want to indicate this parastructure behaviorally, e.g. repelling elements might be actually seen to "fly apart". Such parastructural behavior is currently rather uncommon, but a framework for arbitrary structurality should provide for it. (For a discussion on the subject of "anti-links" see [8].)

HYPERTEXT IN THE OPEN AIR: A SYSTEM-LESS APPROACH TO SPATIAL HYPERTEXT

INTRODUCTION

This paper [1] presents an ongoing project called (loosely) The Frame Stack Project. The term 'frame stack' describes an interface concept which I have been using for a number of years, and have described previously (Rosenberg 2001). For the sake of completeness, the frame stack concept will be reviewed briefly here. A frame stack is a user interface for overlaying word objects on top of one another, while still allowing them to be read legibly. In its starting, or "closed" state, all of the word objects simply appear one on top of the other, with no particular graphical indication of any interface components. The bounds of the frame stack constitute an active mouse-over hot-spot. When the mouse enters the bounding box of a closed frame stack, a stack of frames is displayed, one for each element of the overlay. These frames serve as mouse-over hot-spots to allow navigation of each member of the overlay; note that the frame stack explicitly embeds the peer structure of the overlay. An open frame stack is illustrated in Figure 1. (The same frame stack is shown closed in Figure 2, surrounded by a "morphic halo" (see below). The frame stack concept works well with nested collections, since it allows a frame stack to be opened inside an outer frame stack. One drawback to the frame stack concept is that it is quite expensive in terms of screen real estate. Assuming all frames in one stack are the same size, this must be large enough to accommodate the largest element — with sufficient room for a bit of margin.

Figure 1: An "open" frame stack

In the past, "frame stack" has been more of a conceptual artistic framework than an actual implementation; until recently it would not have been possible to examine any of my finished works and uncover actual objects identifiable as frame stacks. This paper describes the implementation of the frame stack concept as actual working code. It has resulted in a kind of personal spatial hypertext authoring system, but of a somewhat different kind than usual. Rather than an "application", within which spatial hypertext development takes place, the Frame

Stack Project consists of a lightweight set of classes that operates within a generic object desktop. This allows word objects — complete with interactive behavior — to be simply "loose on the desktop" — a concept discussed below as feral structure. Feral structure is closely related to the classical philosophy of spatial hypertext: that the user must have the ability to postpone creation of structure. The specific structural operation which is postponed in the case of a feral object is parenting. Accreting of "parentless" objects is a method of facilitating, in software, a time-honored central part of poetic practice: scrap collecting. Most poets keep some form of notebook, in which scraps are accumulated, often without any clear idea at the time it is first written down what the ultimate destination of the scrap will be. A scrap is thus an inherently parentless object.

The generic object framework within which the Frame Stack Project is realized is called Morphic (Smith et al, 1995) (discussed below), and is provided off the shelf in the programming environment Squeak (Ingalls et al, 1997). Of course it is somewhat disingenuous to describe this approach as "systemless" — after all Squeak could certainly be considered a system. But can it be considered a "hypertext system"? Most researchers would probably agree this would be a stretch.

MORPHIC

Morphic is a user interface paradigm providing a wide variety of graphical facilities. A "morph" is an instance of the class Morph, or one of its subclasses. Morphs can contain other morphs; when a morph is moved on the desktop, its submorphs move with it. When a morph is selected by a mouse-click combined with a keyboard modifier (which depends on the operating system) — e.g. "command-click" on the Macintosh — a set of icons called a "halo" appears. (See Figure 2). The halo mechanism provides an interface to a set of generic behaviors of all morphs, which subclasses are free to override if necessary. Among these behaviors are: moving, pickup, resizing, menu, iconify, and delete.

Figure 2: A closed frame stack surrounded by a "morphic halo"

2.1 Pickup

A morph can be picked up, by dragging the top middle icon from the halo. This means not only moving its position; when the mouse is released at the end of the drag, if the mouse cursor has entered the bounding box of another morph, the morph being picked up is dropped into the target morph, becoming a submorph of the target. The receiving morph may have been set to reject drops; in this case

the picked-up morph will "fly back" to its starting position. It is interesting to contrast this concept with the way that aggregation works in familiar spatial hypertext systems such as VIKI or VKB. VKB, for instance, provides a mechanism called Collections. In order to aggregate spatial objects, they are placed "into" a collection. A bit of text, by contrast, "isn't" a collection; in VKB a text object is a terminal node in the structure, and cannot have subobjects. Morphic, on the other hand, assumes that (1) any object may contain subobjects and (2) the decision whether or not an object should contain subobjects is made "on the fly". (To a poet this means: one can change one's mind about this!) I.e. for a morph to change state in either accepting or rejecting pickup, or to contain or not contain another morph as a submorph, is not a change of class. By contrast, if a phrase in VKB is treated as an "object" (terminal node) and the document author suddenly decides this phrase needs to have subobjects, transforming the phrase so this is possible is a very heavy-weight activity. The user must (1) create a new collection; (2) copy the phrase into the title of the collection, (3) delete the original phrase. The Morphic concept of submorphs may be said to be closer to the spirit of spatial hypertext than a structural concept like collections, in that the decision whether an object is to contain subobjects can be postponed. (As we speak repeatedly in spatial hypertext discussions about postponing the realization of structure, it must be emphasized that at the moment when structure procrastination ends, conversion to structure needs to be as light-weight an activity in the user interface as possible!)

Morphic makes a distinction between a morph accepting drops and a morph containing submorphs. Rejecting drops is a "state property" — i.e. it's determined by the current state of an instance variable, which might change. However all morphs contain the mechanism needed to contain submorphs, and even if a morph rejects drops, submorphs can still be added programmatically by a Smalltalk method. Thus the state of rejecting drops may be considered an interface property, since the rejection behavior is enforced by the interface and can be overridden. It should be noted that Morphic does contain one serious user "hazard": if a morph is enabled for receiving drops (in order to "build it up" interactively), and then that morph becomes "finished", the user may forget to turn off acceptance of drops. If this morph becomes embedded in a larger morph, it may "attract" a drop that the user thought was going into the larger morph. I.e.

in the version of Morphic used for this project there is poor visual identification of the target of a drop.

2.2 Event Handling

An instance of a subclass of Morph can "register" with the Morphic event handling system that it wishes to receive events, such as mouseEnter. In this case the object will be sent a mouseEnter message. This allows any morph to create its own user interface behavior. There is an important point here. Unlike complex paradigms such as the Model View Controller paradigm (Krasner and Pope, 1988), a user interface designer employing Morphic need only subclass at a single point in the class hierarchy. Thus, to create the class FrameStack, it was only necessary to subclass RectangleMorph. An MVC approach would require subclassing at three places: for the model, view and controller. Smalltalk presents the novice with a gordian knot learning problem: where to subclass? Most programmers learn a new language by writing code. But in Smalltalk, one cannot simply "write a program"; all code must go into classes, and one cannot know where one's class should go in the hierarchy without learning the hierarchy, which one can't do without learning Smalltalk ... The genius of Morphic is that it provides a ready set of classes available for subclassing whose function is completely intuitively clear. E.g. it was intuitively obvious that some of the important classes for this project needed to be subclasses of RectangleMorph.

2.3 Menus

A morph inherits a "standard menu" listing generic Morphic behaviors, but the programmer creating a Morph subclass can easily customize this. This provides an easy method to attach an interface to one's own classes. Figure 3 shows the "morph menu" for an instance of the class FrameStack. The entries at the bottom correspond to custom methods. E.g. reframe — the most complex piece of code in this project — tells a FrameStack to abandon its existing frames, make new ones for each of its submorphs, and resize itself appropriately.

Figure 3: The "Morph Menu" for a frame stack

It is also worth commenting in more detail on the menu choice seen in Figure 3 "start playing". When a frame stack is "playing", as the mouse enters its boundary, the frames for the submorphs are made visible, and one of the submorphs is selected at random to be on top. The frames are opaque white, and are active mouseEnter regions. As the mouse enters one of the frames, it and its des-

ignated morph are brought to the front, making them visible. When the mouse leaves the boundary of the frame stack, all of the frames are made invisible. When the frame stack is not playing, it is impervious to mouseEnter. In effect, "playing" constitutes enabling the run-time behavior of the frame stack object. The point here is that unlike environments like Flash, where authoring and run-time behavior are so completely separate that they take place in separate applications (with widely divergent kinds of licensing!) in Morphic a distinction (if needed) between authoring and run-time behavior can take place in a highly dynamic granular fashion as a state property of individual objects. Run-time behavior for an object can be "left on" until it "gets in the way"; at that point it can be turned off for that individual object. This subject will be discussed further below.

STRUCTURE VS. PRESENTATION

It is customary in hypertext that structure and presentation should be separated. Following a link is a structural operation. There may be a wide variety of ways of presenting this operation to the user, even though the structural operation is in each case the same. In spatial hypertext, however, we take a different point of view. Presentation is what "replaces" structure at a point where the user is not willing to commit to structure. Or, to put it somewhat differently, "spatial structure" and presentation are inseparable. Thus for instance while "nearby" is a presentation property, it is of essential importance in replacing a structural concept in spatial hypertext. Morphic provides a ready stock of presentation abstractions which can be used by spatial hypertext objects as a kind of off-the-shelf toolkit. In this section we review some of them.

3.1 "The Front"

Morphic provides a "layering order"; it knows which objects are in front of other objects, and can render them appropriately. A morph can be brought to the front or sent to the back. This operation does the right thing regarding the submorph hierarchy. E.g. sending a morph to "the back" sends it behind all other submorphs of the same parent, but does not send it behind its owner. The presentation concept of bringing a morph to the front can substitute for the structural operation of "navigating to" the morph. For instance, consider a "card interface": to the user it appears that one is navigating among a set of cards, which contain various objects. Some of these act as buttons, which take the user to other cards. Provided that each card is opaque to any cards behind it, this interface can be implemented in Morphic by having each card be a rectangular morph of the same

size and position and then "navigating to" a card by simply bringing it to the front. Technically, the user is "at" all of the cards (at once) but from the point of view of user experience, the only card that is visible is the one on top (at the front, in Morphic terminology).

Using the ability to bring objects to the front as a substitute for structural navigation can have some tricky consequences. Just as an interface can suffer from contention over screen real estate, there can be contention for the front. For instance, when bringing up a halo on a morph, an inconveniently placed frame stack which is playing may bring a frame in front of a halo. (This same effect can be seen on some web pages, where simply moving the mouse — e.g. from a javascript button to the menu bar — can inadvertently trigger a rollover.) Mitigating this contention can require forgoing behaviors. A frame stack which is occluding some interface object which one wishes to be at the front can simply be told to stop playing; since the interface objects are likely to only be present during authoring of some nearby object, the frame stack can be told to resume playing later, with only minor inconvenience. In other cases resolving contention for the front can be more painful. The Frame Stack Project utilizes a package for Squeak called Connectors (Konz) which allows creation of connected diagrams. In this package a "live connector" also makes claims on the front. As of this writing, to resolve this conflict has meant restricting Connectors objects to their own space, forgoing "live connection".

3.2 Visibility

Hiding an object or making it visible is another way that a presentation operation may substitute for a structural navigation. Again, Morphic provides the ability to hide or show any morph. It is surprising that hiding of objects is not an operation commonly supported by spatial hypertext systems.

IMPLEMENTATION SPECIFICS

The hierarchy of classes created for the Frame Stack Project is shown in Figure 4. The amount of code in these classes is so small that it is an almost infinitesimal fraction of the corpus of Squeak / Morphic. I offer this not as an apology, but rather as testimony to what can be accomplished by an individual cybertext author using the "open-air subclassing" approach on top of a rich generic object framework like Morphic. The Frame Stack interface is particular to my own artistic practice; other writers will have drastically different needs. It is unlikely that very many cybertext authors will find the Frame Stack Project code directly use-

ful; I am offering it more as a kind of living example of what can be accomplished using this method.

(Array)
FrameStackRectArray

(Form)
FrameStackGlyphs

(Object)
FrameStackExport

(PasteUpMorph)
FrameStackCard
FrameStackRelation

(Rectangle)
FrameStackAttachment

(RectangleMorph)
FrameStack
FrameStackCardButton
FrameStackFrame
FrameStackNode
FrameStackProjectButton
FrameStackQuitButton
FrameStackRectangle

(SketchMorph)
FrameStackSketch
FrameStackScope
FrameStackThumbnail

(TextMorph)
FrameStackText

Figure 4: Class hierarchy created for the Frame Stack Project. Classes shown in parentheses are off-the-shelf classes provided with Squeak.

In the following section, some of the key classes implemented in this project will be described.

4.1 FrameStack, FrameStackFrame

FrameStack is the "signature class" from which this project takes its name. A frame stack is an object with a rectangular boundary and a collection of submorphs for which the frame stack acts as interface. The goal is to provide an intuitive interface by which transparent word objects can be overlaid in the same space — which would normally render them illegible — and allow individual objects to be read by a set of opaque "frames" that are controlled by mouseEnter hot-spots; each frame corresponds to one of the word objects. These frames are implemented by a class called FrameStackFrame. The submorphs — called "exhibits" — for which the frame stack acts as an interface are not "specially" designated in any way; a frame stack identifies its exhibits as any submorph which is not a FrameStackFrame. Thus a new submorph can be added using any means supported by Morphic (such as Pickup, described above) without requiring any special code in FrameStack. I.e. "authoring" a frame stack is as simple as creating a new empty FrameStack (using the Squeak desktop "new morph" menu entry), turning on "accept drops" in the FrameStack, and then dropping morphs into it. Any graphical object can be an exhibit, as long as it is a morph.

In the current implementation, a frame stack is "reformatted" for a change to its submorph population by an explicit reframe method. (Future versions should probably do a reframe automatically in response to various relevant events.) The reframe method discards any existing FrameStackFrame submorphs and then recreates them, sizing them to their designated submorph; the boundary of the frame stack itself is also resized.

In addition to controlling whether a frame stack is "playing" or not (discussed above), another behavior implemented by FrameStack is a "freeze". Normally when a frame stack is told to stop playing, it will be in a "closed" state. (All frame stack frames are invisible, so that all of the exhibits are visible and appear overlaid.) If one of the exhibits needs to be edited, having the frame stack continue to play will interfere, but the "layer" with the given exhibit needs to be "open" so the submorph is easily accessible for editing. Because the mouse is already used to "navigate" among the exhibits of a frame stack, the keyboard is used to

register a freeze. When initially created, a frame stack does not accept keystrokes, but it can be told to do so. Once accepting keystrokes, when a frame stack is sent the 'f' key from the keyboard, it freezes in its current state. This allows the submorph for the layer showing to be edited in place. (Typically editing occurs using off-the-shelf behavior of Morphic; e.g. if the object in question is text, it may be edited using customary text editing mouse moves and keystrokes — as implemented by the off-the-shelf Morphic class TextMorph.) This is consistent with a deeply held philosophy of this project, that the distinction between authoring vs. run-time behavior should be a state property of individual objects, not of "the system" or "environment" as a whole.

4.2 FrameStackRectangle

This class is used to implement grouping. An instance of this class is a transparent rectangular area with various submorphs; it is needed as a specific class mainly to allow grouping in such a way that the mouse events are properly passed through to any frame stack submorphs. Note that this type of grouping is entirely"author-convenience" grouping: the group allows its submorphs to be moved as a unit, including all their behaviors, but the group is not visible to the reader. While such a group is arguably "a structure", it may be described as a purely private structure, much like the corner braces for a painting's stretcher which are invisible to the viewer. Where groups are nested, there may be many different ways such a group may be constructed; e.g. elements a, b, c, d may be grouped as {a, b, c, d} or {{a, b}, {c, d}} or {{a, b, c}, d}; since group substructure is not visible at reading time, the choice among these different substructures may not matter as long as ultimately all the right elements are "together". Note that this position on structure is in total contrast with the point of view of Structural Computing (Nürnberg et al, 1997) where structure is considered primary and structures are first-class objects.

Issues pertaining to grouping in the context of spatial hypertext are discussed in more detail in (Rosenberg, 2004).

4.3 FrameStackCard

The actual cybertexts so far realized in this environment have used an "outer interface" extremely similar to the original card interface provided by HyperCard (Goodman, 1997). This interface assumes a non-scrolling fixed "portal" which does not move on screen; as the reader moves through the piece the content of this portal is changed. In the Frame Stack Project there is no formal con-

cept of portal. Rather, its appearance is created by the cybertext author creating a set of frame stack cards which all have the same size and position on the screen. This class implements a parent-child relationship among frame stack cards, using methods seekParent and acceptChild. When a frame stack card receives acceptChild, a button is created that will bring the child to the front when clicked; the button is a thumbnail image of the child. (At the child there is a method that will set the magnification scale for creating this thumbnail.) An "up-button" is created on the child, that when clicked will bring the parent to the front.

Consistent with another major philosophy of this project, a frame stack card may have no parent. (There is also a method of FrameStackCard called unparent, which will delete the parent relationship and render a card parentless.)

FrameStackCard is a subclass of an important Morphic class called PasteUpMorph, also known as a playfield. This class is the basic form of Morphic "canvas", and provides many facilities for graphical editing. The Squeak desktop (known as a "world") is in fact a playfield.

4.4 FrameStackRelation, FrameStackNode

Relations as a built-in structural concept were first seen in hypertext systems in Aquanet (Marshall et al. 1991) but have seldom been implemented since. In addition to the kind of direct juxtapositions realized as frame stacks, my artistic work for many years has utilized a diagram notation which has an explicitly relational model. Fragments of this notation can be seen in Figure 5. In Aquanet, a relation is primarily a structural concept, and secondarily a relation may have a graphic content by which the relation is rendered visually. In the Frame Stack Project this relationship is reversed: the relation is primarily a graphic object, that secondarily has structural properties. The objects that are "tied together" by the relation are pointed to by FrameStackNode objects as an attachment of the frame stack node. FrameStackNode has an attachToNearest method which attempts to locate the nearest suitable target; FrameStackRelation has an attachAll method that executes attachToNearest on all nodes. This is the closest thing there is in the Frame Stack Project to a spatial parser.

At this writing, the frame stack relation concept may be described as "weak structure". The original implementation of the Frame Stack Project (in which several artistic works were completed) did not have FrameStackRelation implemented. Instead, in these works, relational structure was denoted by entirely visual means. (These works were later revised so that all of the PasteUpMorphs consti-

tuting "visual relations" were replaced by actual FrameStackRelations.) Even now, the structural concept is still weak in that (1) visual structure and node / attachment structure might be inconsistent; (2) the node / attachment structure has no consequence for the reader. Errors in the node-attachment structure do have consequences for the ability to export and to "walk" the structure (see below).

4.5 FrameStackExport

FrameStackExport's main function is to be a "state holder" to support export of information about a project to an external file format. Currently the only format implemented is text; although all structural information and text content is exported, spatial information, relation geometry, and text metadata (e.g. font) are not yet exported. Because the relational structure may contain loops, the export algorithm, though not difficult, requires more than one pass. This class includes a "descent" method whose argument is a Smalltalk block, allowing it to serve as a kind of open framework for code that needs to operate on all objects in a project. "Ad hoc" use of this method to walk all objects in a project proved very useful lately getting around an obstacle in porting the Frame Stack Project from Squeak 3.6 to Squeak 3.8. In general, such a method is enormously valuable for cybertext preservation, and provides a kind of generalized "extensibility hook". It is unfortunately true that our record in spatial hypertext of providing extensibility is not good.

4.6 Fonts

Fonts are an extremely tricky issue in any discussion of cybertext authoring systems. It is customary among hypertext system designers to assume that fonts are someone else's problem; e.g. the native operating system windowing system is presumed to provide fonts, the user may have fonts of her own, etc. Scalable outline fonts, such as TrueType or Postscript Type 1 fonts, are a form of intellectual property subject to their own system of rights. A cybertext author wanting to control the exact appearance of the text is thus confronted with a difficult dilemma: embedding fonts in a cybertext may create unpleasant rights problems for distributing the cybertext. Technologies like Flash seem to allow distribution of cybertexts with embedded fonts in ways that have apparently avoided this problem, but at the cost of a heavy-weight distinction between the authoring environment and the run-time environment. An important goal of the Frame Stack Project was to be able to support creation of cybertexts with embedded fonts that the cybertext author can edit. As of the time this project began, the native font

system of Squeak is bitmapped. The decision of whether to use bitmapped rather than antialiased fonts was one of the more aesthetically difficult decisions made during this project. In the end, a set of fonts was created based on outline fonts believed to be unencumbered; from these, screen renderings were imported into a Squeak font editor to create bitmapped fonts with a close aesthetic resemblance to the effect of antialiased fonts on screen. The FrameStackGlyphs class was created to allow importing an entire character set at once. As the Squeak font system evolves, the fonts used will probably change.

FERAL STRUCTURE

[2]. There is a great deal of research involving integration of hypertext systems with a larger computing environment, particularly in the OHS community. Hypertext has certainly had a wider perspective than just "the confines" of hypertext applications for quite a long time. Still, it is most common for hypertext objects to be found inside hypertext systems. While the Squeak desktop is not the native operating system desktop — though it could become the native OS desktop; see (Smith et al, 2002) — it is certainly a "generic object desktop" in which the user could spend the entirety of her time and which is not especially devoted to hypertext. The desktop is the cyberspace equivalent of the open air. A desktop such as the Squeak World allows objects to be simply "loose" in the open air, much as a physical desktop allows physical objects to be loose on its surface, without being placed in a drawer. The appeal of such freedom is similar to the attractiveness of spatial hypertext itself. Among the features offered by feral structure are:

Objects near at hand are presumably prioritized.

A disposition of the object can be postponed.

A persistent desktop allows work to be resumed in exactly the state it was left in a previous session.

It is particularly important to note that feral structure is ideally suited to collecting cybertextual scraps where the destination of the scrap is not known at the time it was collected. As mentioned above, there is a deep historical affinity for poets in particular to write by a method that in part involves accumulating materials in notebooks. Systems such as Flash, with their extremely heavyweight distinction between authoring and run-time, raise profound difficulties for collecting cybertextual scraps. Figure 5 shows a screen dump of the actual live Squeak desktop for my work in progress as of July, 2003. Note there are several

objects placed on the desktop wherever I found it convenient to work with them: some are frame stacks or frame stack rectangles, some are frame stack texts. Note the objects in the top left corner. These are iconified morphs. The ones marked "playfield" are frame stack cards which are more or less finished, but have not yet been integrated into any higher level of structure.

Figure 5: Screen shot of live work in progress, July 2003, showing objects "loose on the desktop"

Long time users of (say) VKB may wonder why there is any difference between the concept of feral structure as articulated here and the VKB "root collection". After all, in VKB no one is obliged to make collections; one may place all of one's objects in the root collection. I.e. VKB allows a structure which is "flat". What is the difference between a flat structure and feral structure? Perhaps one could argue that this distinction is simple hair splitting, but the major difference is that an application like VKB is not a generic object system, in which any kind of object (with any kind of behavior!) can be placed. A VKB collection can only contain the kinds of objects that have been specifically implemented in VKB. It is not "the open air", but rather a very special atmosphere in which only a severely limited variety of creatures can breathe. While it would be easy to imagine the Squeak desktop as the native OS desktop, this would not be possible with VKB without a very significant amount of work.

5.1 Structure Objects vs. Activity Results

In speaking about feral structure, it is useful to consider what the objects are "feral from". Adherents of structural computing tend to see structure as primary; from this vantage point what one sees in spatial hypertext is often described as "spatial structure". While this approach is interesting and can give some fruitful results, it also misses the essence of spatial hypertext. Consider the word objects in the top right of Figure 5. Why is this cluster in the top right? Is its being in the top right really a case of "spatial structure"? In this case the author is available to testify. I often tend to put a cluster in the top right of my workspace when (1) it is available to use as material in a larger structure; (2) I don't want to forget about it; (3) I have moved it out of the way to clear space nearer the center of the workspace for the objects I am actively composing at the moment. I.e. the location of the cluster in the top right is the result of an activity. Perhaps its spatial relation-

ships in its out of the way location can be fruitfully analyzed in structural terms, but what mattered at the time the object was moved was: getting it out of the way. Any concern about structure was at best secondary. Structure is important, to be sure; but even more important is the workspace as an arena for word object activities. Certainly in this arena structure happens. But also structure "unhappens". Word object activities occur that may or may not give structural results.

What open-air word objects may in fact be "feral from" is exactly: structure! A generic desktop which allows interactive word objects to be loose on the desktop allows such objects to simply be present. This allows the author to look at them and say: "Hmm. I don't know ... ". The structural relationship of such an object to those already present in the space may be far less important than the simple fact of its arrival.

5.2 Is the Native OS Desktop a Spatial Hypertext?

It can be argued that a desktop is not truly the computer equivalent of the open air unless that desktop is the ultimate "native" operating system desktop. That raises an interesting question: should we consider native OS desktops as "already" spatial hypertext systems? Many users certainly place a great deal of information on their desktops, and some users become completely lost if a desktop icon goes missing: they navigate not through the file system, but spatially on the desktop.

There are no commercial operating systems that have a desktop with the object power of anything like even a fragment of Morphic. Perhaps we can look forward to this in the future.

REFLECTIONS ON USAGE

The concept of usability takes on an odd cast in the context of a personal authoring system. How should the author of a personal authoring system carry out an unbiased usability study? This is clearly impossible.

It will have to suffice for me to simply offer anecdotal evidence. Based on approximately three years of creating finished works in the Frame Stack Project, I can say that total elapsed time to complete such a work is cut by a factor of about 3 from my previous methods. More importantly, (and even more anecdotally, alas) the feeling of composing in this environment is substantially different than it was using tools like HyperCard. When writing in the Frame Stack Project, the word object is a true object, and can easily become a finished "interactive scrap" during a single session. By contrast, using previous methods the objecthood of what

appears on the screen as a word object is a mere facade; inside the work there is no real object, and it might have taken weeks after all aesthetic decisions were made before there was any interactivity present at all. Writing by such methods relies on a completely non-interactive document which Bootz (Bootz 1997) calls the texte auteur, which provides a kind of implementation specification for how a cybertext is to be assembled. While opinions can differ concerning what the term "interactive writing" might mean, it is hard to call a writing process interactive if interactivity appears only at the end of a long process, taking weeks or months in which there are no interactive objects present.

The intensely granular individual object nature of the distinction between authoring and run-time achieved in the Frame Stack Project simply gives a different feeling to the act of writing. It allows interactive writing in the true sense of the word. The act of writing involves countless acts of reading by an author of pieces of the work in progress; the writing process intimately involves intense feedback from reading material already written. For writing interactive works, having fragments of objects already composed "loose" and at hand greatly facilitates this feedback. Many effects of reading can easily persist into the writing process; such persistence is inhibited by the need to transition back and forth from separate environments for authoring and run-time. (See (Rosenberg 2004) for further discussion of persistence effects.)

FUTURE WORK

Because the Frame Stack Project takes place on the substrate of an open spatial object framework, the way is wide open to add new kinds of objects. My current poetic work includes some new object types for spatialized linear word objects. There is much more work to be done in export functionality; ideally enough should be exportable to allow complete reconstruction of an artistic work. Some form of XML as an export format would clearly be highly desirable. Another useful export format would be Smalltalk code; using such an export, a project could be reconstructed by the simple expedient of filing it in.

The FrameStackRelation concept currently has no built in support for "off-card" nodes. In fact the relation model so far implemented has no associated behavior whatsoever. It should be easy to associate any objects in a project into a relation; there is much work to be done in this project to support "distant" relations.

8 ACKNOWLEDGMENTS

This work would not have been possible without the immense contributions of many other people, specifically: Allan Kay et al for Squeak, John Maloney for bringing Morphic to Squeak, Ned Konz for Connectors, Boris Gaertner for his font editor, and of course Cathy Marshall, Frank Shipman, and many others for the concept of spatial hypertext in the first place.

NOTES

1. This paper is a revised version of (Rosenberg 2003).

2. Walker (Walker 2005) presents a somewhat different concept of “feral”; in her concept objects are feral in the Internet at large. As she cites in her paper, the discussion here — as originally found in in (Rosenberg 2003) — predates her discussion.

REFLECTIONS ON SPATIAL WRITING IN PLACE

INTRODUCTION

Spatial hypertext is often used to create documents whose ultimate destination is some other medium. E.g. VKB [7] makes an excellent brainstorming tool, and ART [8] was explicitly designed as a spatial editing environment for creating linear texts, which would presumably be consumed using conventional methods. Elsewhere [6] I have described the Frame Stack Project: a light-weight set of classes for the programming environment Squeak [1] which provides me with an authoring environment for interactive word objects where any distinction between "authoring" and "run-time" environments is a granular state property of individual objects. In this authoring environment, word objects "play in place" without having to "execute" a separate browser or run-time player. This paper, rather than focusing on the details of that environment, reflects on the experience of using that environment over a period of about a year and a half to actually write a number of interactive poems, and contrasts this experience with earlier experiences using other kinds of environments.

A caveat: The Frame Stack Project, by design, has been aimed at creating a personal authoring environment. It reflects requirements of my own rather peculiar artistic practice, and reflected my inability to meet those requirements with the off-the-shelf multimedia authoring environments I had tried. This paper is even more personal — exploring my own personal experience at using this intentionally personal authoring system. To those accustomed to usability studies, the value of such results on a sample of one may seem inherently dubious. To this I can only reply that personal empowerment has always been an important goal of tool builders, and at a time when toolkits for creating interactive word objects that play in place are not common, someone has to go first and actually write in such an environment. A document may have many readers, but typically has only a small number of authors. The authoring experience is — almost by definition — an intensely personal one.

GRANULARITY

At what scale does writing begin? One can ask this question in general, or with respect to a particular session. If the question is directed to a particular session, then presumably there are some prior word objects that have been created, so the granularity will be influenced by what has been left "pending". Writers vary enormously in how they actually write. Historically, poets have had a longstanding fondness for working at least in part bottom-up: the poet may "hear" a line, which is collected as a scrap and placed in a notebook; several such scraps may then come together into a larger unit. As a poem proceeds, the granularity may shift back and forth, from extremely small-scale units: the word, or the line, to larger units, such as a stanza, and ultimately to the poem as a whole treated globally.

It is at the smallest level of granularity that an environment enabling spatial writing in place can make the most dramatic difference. As the amount of text in a word object diminishes, the "mechanism" of that object tends to loom at a larger scale with respect to the text itself. Consider a simple spatial cluster of multiple elements. Prior to using spatial writing methods, I might have indicated this in the act of composition by a kind of intermediate notation, so that a cluster with phrases 'life splay ratchet' and 'sentience tarp' inserted into the phrase 'to name blank escrow __ parole' might have been indicated this way:

to name-blank escrow <life splay ratchet : sentience tarp> parole

Of course this intermediate notation has nothing about it which is spatial. (And certainly nothing which is even slightly interactive!) Even worse, at small-scale granularity, the "mechanism" of the intermediate notation (< ... : ... >) tends to become extremely obtrusive with respect to the words themselves.

In an intermediate step, I wrote some pieces using VKB as a prototyping environment. (VKB was not suitable as the final destination environment for any of my pieces, because it does not support the kind of behaviors my work required — specifically the frame stack behavior. See [5] for a description of this behavior.) Figure 1 shows how a prototype of this phrase might have been rendered as a "mockup" in VKB. Figure 2 shows a screen capture of the actual interactive fragment as composed "in place" in the Frame Stack Project environment.

There are several notable differences between the approaches shown in the two figures. Even though the approach shown in Figure 1 is to some degree spatial it is not fully spatial. This is because the "final" spatial relationships of the word elements in the word object being assembled are not there in the prototype. VKB does allow transparent text, and does allow objects to be placed on top of one another, but doing this can make the words illegible, and there is no user interface behavior to navigate among the layers. (It is exactly this problem — navigation to achieve legibility among objects placed on top of one another — that the frame stack concept is designed to achieve.) Of course, it goes without saying that the interactive behaviors present in the final word object are not there in the VKB prototype.

What has become apparent to me after several months of writing spatially in place in the Frame Stack Project environment — which was certainly not obvious to me when I was writing using the prototyping or intermediate notations methods above — is that these prototyping environments induce a bias in how I would write. Because the prototyping mechanism itself becomes so much more obtrusive at small-scale granularity, these prototyping methods tended to skew the content of what I would write to a larger scale granularity, in which there were simply more words in the leaf-nodes of the structure. I certainly did have very granular objects from time to time in my earlier writing — e.g. clusters with one word per layer — but they tended to be rare.

GROUPING

Grouping is an important function present in nearly all graphics software. Historically, spatial hypertext has had a tendency to take a different approach to grouping than your typical drawing program. Grouping in commercial graphics software tends to have the following characteristics:

- Objects are grouped interactively by muti-selection followed by the action of grouping.

- There is typically no distinctive user interface for a group.

When a group is selected, handles appear around the entire group, and menu selections pertaining to groups may become enabled. The appearance of these handles is the same as the appearance of handles on "atomic objects" like a line or a rectangle; when the group is not selected its "objecthood" is invisible.

By contrast, grouping in spatial hypertext sometimes relies on spatial parsing, which is "automatic" rather than interactive.

For artistic work, such as poetry, the technical issues pertaining to grouping take place against the following backdrop:

- Structural groupings may be a natural part of the work itself.

Historically, poetry has involved such units as lines and stanzas; the units of interactive electronic poetry may be different, but may still be present on a range of scales.

- Of immense artistic importance is the concept of whether a unit is finished.

As composition proceeds, units may exhibit a great variety in their degree of being finished. Artists frequently change their minds about whether a unit is finished. When a unit is not finished, an additional piece may be a candidate for inclusion in the unit. This means that at some stage in the composition process, membership of the additional piece in the unit may be uncertain or ambiguous. Of course, it exactly to be able to model such ambiguities that is one of the prime purposes for which spatial hypertext was created in the first place.

In my own practice I have found it most useful to make a group where: (1) a unit is probably finished; (2) I am almost certain to want to move the components together, not separately. I.e. the group is a set of spatial elements where I want the internal spatial relationships of the elements to persist. If I am thinking about adding a new element to the unit, I would tend to place it "nearby". It is important to note that as composition proceeds, a unit that is not finished may be drastically not finished. For instance, a phrase may have holes in which the

words haven't been chosen yet. (In my own practice I tend to "write" these holes explicitly, often with an indication of what kind of prosody they should have, but without the specific words.) Of course a unit which is "drastically unfinished" may actually end up being discarded (or cannibalized for its parts!). In this situation I am personally very unlikely to want to include in a group a unit which is drastically unfinished.

I.e. the act of grouping is itself an artistic act, with artistic consequences. This has direct ramifications for spatial parsing. If an algorithm is forming groups — virtual or otherwise — on behalf of the author, that algorithm has artistic consequences. It could happen, of course, that the author of the spatial hypertext system has happily chosen an algorithm that is exactly in concert with the aesthetic of the document author, and the document author is getting just what she wants from the algorithm. But, realistically, the opposite might also happen. This suggests that if grouping is done by a spatial parser, at a minimum, tuning parameters for how that algorithm works must be accessible to document authors, and ideally the entire algorithm itself should be accessible to allow modification. Historically, spatial hypertext systems have tended not to be open source. Where algorithms have artistic consequences but the source is closed, hypertext system authors are giving document authors only the option of turning the algorithm off.

One possible compromise where open sourcing the entire spatial hypertext system is not an option is a form of plug-in architecture. If a spatial parser is implemented as a plug-in, with a clearly documented interface, in theory at least document authors who are unsatisfied with the off-the-shelf spatial parser could write their own. It might also be possible for just a spatial parser plug-in to be open sourced, even though the entire spatial hypertext system for one reason or another cannot be open sourced.

History vs. Structure

An interesting issue for grouping concerns the history of how the group is constructed. Consider a group made of three elements, a,b,c. Now a fourth element, d is constructed, and the document author decides that all four elements need to be grouped. There are two ways to do this: (1) the element d can be added to the group {a,b,c}. (This may require a, b, and c to be ungrouped first, and then

reselected, which can sometimes be cumbersome. Or perhaps d can simply be dragged into the group. This case would require a "heavy-weight visibility" for the group boundary — which may have undesired aesthetic consequences.) (2) A new group can be created as {{a,b,c},d}. This is likely to be much faster than method 1, but can cause some surprises if the document author needs to edit the group later and forgets the history of how it was created. Of course these are two drastically different structures! The question is: does it matter?

There is a strong tendency in the hypertext community toward heavy-weight attention to structure, even to the point of placing structure at the center of the universe (Structural Computing [2].) For those from this point of view, the idea that such a drastic structural difference "may not matter" is likely to sound like heresy. But consider the case where (1) grouping is done for the convenience of the document author; (2) the purpose of grouping is simply to preserve internal spatial relationships among the elements; (3) the group as a structure is not visible to the document reader. In this situation, the document reader cannot distinguish the two cases above. If the two cases appear identical to the reader, then we are left with simple author convenience to determine which case should apply. It should be noted that the kind of grouping typically found in graphics software does tend to have all three of these properties.

However: some spatial hypertext systems seem to make the assumption that structure always matters. E.g. in VKB, it is difficult to group objects except by making them members of a collection. The boundaries of a collection cannot be made invisible. VKB collections are not well suited to implementing structural ambiguity.

FEEDBACK

All writers know that the writing environment influences what is written; writers tend to be very particular about their work circumstances. In the days before personal computing, that might mean a very specific kind of pen, an exact brand of notebook, a particular individual typewriter, etc. We have barely begun to explore how this process works in "new media" forms of electronic writing. As a writer works, the parts of the composition already present exert a power-

ful feedback on what gets written next. Particularly in artistic or literary writing done with a strong "bottom-up" methodology, an immensely important part of the composition process is deciding whether a possible "extension" of an element already composed fits. (An extremely significant strain of poetics prominent in American poetry in the latter half of the 20th century, and highly influential on a whole generation of poets, was known as Projective Verse [3]. A major component of this poetics may be described as composition by induction. (Robert Duncan described this when speaking about poetry using Olson's term Composition by Field.) The rough idea is that if a continuation occurs to the poet which works in a localized area, that continuation should be accepted, without having to match it against a "global scheme". Thus Projective Verse placed great importance on the question of whether "that next bit" — induction — fits.) When writing using either a prototyping environment or an authoring environment which is distinct from the run-time environment, the writer needs to use a kind of extrapolation process to imagine the result from "inside" the writing environment. To understand the implications of having to extrapolate — which requirement is of course absent from an environment that supports spatial writing in place — I will next explore some specific issues related to this extrapolation process.

Time Sequence Choices

Hypertext (and spatial hypertext is no exception in this regard!) tends to place interactive devices amid the words. Even if the structure used is inherently conjunctive [4], there will be some history of use in operating these devices: the user has a particular time sequence of having made various "moves" operating the interface. In the case of a cluster where the elements are meant to be peers, there will nevertheless be some element which in a particular case was visited first, and some which was visited last. When using an environment that allows writing in place, the time sequence in question as experienced by the writer simply is exactly the time sequence experienced by the writer acting as reader. Where the word objects do not play in place, however, this time sequence — along with other behaviors — must be extrapolated. How many such time sequences has the author extrapolated? In the case where elements are not reordered dynamically, the document author is likely to have a "favored order" in which elements are extrapolated. The feedback effect of "playing the objects" in a different order will not be experienced. Of course the document author can obtain this feedback

by interrupting the composition process, in the case where "playback" does not occur in place, but this interruption may be disruptive. The feedback may occur, but may not be "invited".

Persistence of State

When an interactive word object is played in place, it is left in some particular state. That state is likely to be a state that persists from state that was dynamic while the object was playing. By contrast, for an authoring system where the objects do not play in place, state within the authoring system is simply not dynamic at all, but static. Persistence of state is thus another aspect that a document author must extrapolate when trying to compose with an authoring system where objects do not play in place. Consider a multiplicity where one element "drowns out the others" — an effect that the document author does not desire. This may only be apparent when the word object is played. If feedback during the composition process is used to compose further elements, this drowning out effect may not have played out. A compositional continuation may have been written based on a voice which "did not sing as loud" as the voice during object playback.

My own practice has for many years involved putting words on top of one another while allowing the layers to be read legibly — an activity to which I am deeply committed. When the resulting cluster is "closed", there will be a lot of variation in the legibility of individual words. Some will be obscured completely, while others will be completely legible. In some cases all the words in the cluster can be read easily once the reader has read them in individual layers; in other cases almost nothing can be read, but an occasional word may "stick out". This is a persistence of state: when the object is closed, it stays closed until it is activated interactively. The varying weight given to the words by persistence of state is only observed once the word object has been assembled. It could be argued that there is no reason to have the word object play in place to be able to observe this: just the graphical assemblage displays persistence. The answer, of course, is that there is a vast difference between state that persists after the experience of a dynamic state, and a completely static state. The word that "sticks out" from a phrase you have read will have a different resonance from the same instance where the context of the entire layer has not been experienced.

It goes without saying that such persistent state effects are nearly impossible to extrapolate when a work is composed by a prototyping process where the ultimate graphical look of the word object is simply not available until "later". This is a form of extrapolation that is nearly impossible to do.

Persistence of Aesthetic Effect

When a word object is "played" (or perhaps "operated" is a better term) there is an aesthetic effect. How long does this last? How resistant is it to "noise" — in particular the noise of having to transition from an authoring environment to a run-time environment and back? If the aesthetic effect of playing a word object in a separate run-time environment is significantly dissipated in the transition back to the authoring environment, this is likely to create no harm for the ability of the author to evaluate the word object — in effect in isolation. But it may severely inhibit the ability of that word object to exert feedback on the composition process.

It is difficult to describe in words, but there is a very strong aesthetic feedback effect from graphically manipulating a word object that you have just played. Consider the case of a mostly finished cluster, and a possible candidate cluster which the document author may wish to include. These two objects can be played, dragged close to one another, then played again in a nearly seamless act where the time it takes to spatially "indicate candidacy" is within the persistence of aesthetic effect. This provides the author with a strong preview of what it will be like to combine the two objects. It is important to emphasize that such "compositional moves" may yield failure. From the persistence of aesthetic effect, the author might decide that the combination being considered simply doesn't work. It can save a document author considerable disruption to be able to reach such a judgment without going through the heavy-weight process of creating and then evaluating a combined object and then deciding the result was not successful and should be undone.

Software designers typically believe that their job is to create systems that are engineered for success. But in the arts, creative success only happens amid many failures that are discarded. Thus software which is designed to support artistic endeavors must also be engineered to support failure. Spatial hypertext

has long been associated with such concepts as ambiguous structure, emergent (or incipient) structure, and other structural concepts where structuring is not "completely done." The result of such incompletion may go both ways: a structure may emerge, or the document author may decide that structure among a set of elements has failed, and the set is dissolved. Spatial hypertext research needs to devote as much energy to understanding and supporting the process of structure failure as it does to structure emergence.

Cognitive Load

This is a very difficult issue — the elephant at the dinner table for cybertext authors. Is the cognitive load that the word object imposes on the reader manageable? Of course a document author can conduct usability studies to try to answer this question, but a usability study is hardly feasible as part of the composition process at a granular level. As a first resort, a document author is most likely to just play the object and ask herself: is this too hard. It is quite aesthetically risky to do this by extrapolation. But, document authors will be tempted to extrapolate to avoid compositional disruption if objects do not play in place. How often will the document author click "play in browser"? Every few words? Play the whole piece every few words? Not likely. This has an intimate relationship with the subject of granularity. When the word object does not play in place, the document author's evaluation of cognitive load is likely to be made on large-scale granularities. This means that small-scale pieces that have too high a cognitive load in a context where overall cognitive load is manageable might be missed. Another effect is that evaluation of cognitive load is just less available as part of the compositional feedback process.

CONDITIONAL SPATIALITY

INTRODUCTION

Many forms of hypertext have attached conditionality to various structural elements or forms of behavior. Perhaps the most notable example is the StorySpace abstraction called Guard Fields [3], in which the availability of a link can be made conditional on whether some other node has already been visited. Another example is the Connection Kit [5], an authoring framework for conditional links using JavaScript. However, use of conditionality in spatial hypertext is at this writing unusual. For instance both VIKI [7] and VKB [10] allow an object to be a member of a collection; however this membership is absolute: an object either is or is not a member of a collection. Spatial position also tends to be absolute: an object stays where the document author put it, and its location is only conditional on an author's (presumably persistent) choices, made interactively.

There is no reason why conditionality should not play just as strong a role in spatial hypertext as other forms of hypertext. Introducing conditionality affects many aspects of spatial hypertext; among them are attention and emergence of structure. Conditionality can be achieved in a variety of ways, from simple layering to more complex automatic rules.

Consideration of conditionality in spatial hypertext is still in its infancy, and this paper will only begin to raise some of the issues involved.

STRUCTURE VS. ATTENTION

Management of attention is an integral aspect of spatial hypertext. By placing objects near one another, the document author may be organizing materials so that attention is paid to certain items together. Analogous to the more formal relationships of argumentation, a spatial hypertext author may be creating a relationship of the form "Don't forget A when considering B." Thus an author may be willing to commit to a relationship of co-attention when a more "struc-

tured" relationship has not yet been decided. How this plays out with respect to conditionality has yet to be investigated.

Conditional spatiality introduces some odd paradoxes. If the spatial attributes of an object are not fixed, as they change this will introduce an event into the spatial hypertext of a kind that does not exist in "unconditional" spatial hypertext. This event is likely to mark the affected object, and perhaps draw attention to it in a way that does not occur for elements with no conditionality. For example, consider a concept of layers, and consider two objects, A and B. A is present in all layers, but B is present in some layers but not others. As conditionality triggers a layer in which B is present, B will seem to appear; when this conditionality is no longer in effect, B will seem to disappear. A, however, will remain visible constantly. These events will continue to draw attention to B, making it appear more "prominent". On the other hand, since A is always present, its presence is "stronger". Thus the paradox: while A has "stronger presence" for lacking conditionality, B has "more prominence" by being associated with events that grab attention.

Our nervous systems are strongly conditioned to pay attention to things that move, and to things that suddenly appear. It is likely this is biological; throughout most of evolutionary history, something that moves has a distinct likelihood of being either a potential meal or something that can make of oneself a potential meal. Thus there are important biological reasons for paying attention to things that move. This "hard wiring" of the brain to favor motion is being exploited in web design, unfortunately, as more and more advertisers use animations to forcibly wrest the reader's attention from the "real" content of web sites that is the reason for the reader's visit, making the advertising parasitic on the "real" content — notwithstanding that it is the advertising that makes the site possible by paying for it. Thus issues of conditional spatiality are being played out against a backdrop of contention for attention, which may complicate the analysis. (For further discussion on this point see [9].)

The phenomenon that an ephemeral item may receive attention it doesn't "deserve" due to its sudden appearance may be described as The Flash in the Pan Effect. How to deal with such effects while harmonizing with the objectives of

spatial hypertext document authors and readers is one of the great challenges for user interface design in spatial hypertext.

CONDITIONALITY THROUGH LAYERING

Layering is a familiar user interface paradigm from object-oriented drawing programs. Support for layering in spatial hypertext systems is currently somewhat equivocal. Rather than being an explicit part of the interface, as is typical in drawing programs, layering in systems like VKB tends to be the result of how the user has selected objects. The system maintains an internal concept of layering, which it needs to be able to render objects. The user interface for a typical drawing program allows a layer to be hidden or shown; when a layer is hidden, all objects in that layer remain in the drawing but become invisible. Multimedia systems like Flash [6] also allow objects to be hidden and shown, and this behavior is controllable by means of scripting.

Figure 1 shows an example of conditional spatiality through layering in an unpublished work in progress of my own currently called Diagrams Series 6. 1a shows the simultaneity closed; note the central word cluster appears black and the outer clusters appear gray. This difference is meant to indicate two things to the reader. The central cluster is "active" as a user interface element, while the outer clusters are not active: the region of the central cluster is divided into invisible rectangles which are "on mouseOver" hot-spots activating the layers, whereas moving the mouse over the outer clusters has no user interface effect. The greater "solidity" of "ink" in the central cluster is also meant to indicate that there is no conditionality to layer membership among the phrases in that cluster: each phrase belongs to exactly one layer, whereas in the outer cluster some phrases may appear in more than one layer. Thus the outer phrases may or may not be present, depending on what layer is active. (This is a somewhat weak form of conditionality; what the user sees is conditional based on what the user does, but there is no conditionality to the underlying structure, in the sense of "conditional layer membership".) This piece exhibits some of the paradoxical effects already discussed where persistence contrasts with appearance and disappearance; conditional presence in some layers but not others contrasts with appearance unconditionally but only in a single layer.

GENERIC SPATIALITY

The term "generic spatiality" is meant to provide an analogue to the more familiar generic link [4]; by generic spatiality is meant a mechanism of providing spatial information in response to an algorithm, whose parameters are given in lieu of explicit spatial parameters. An example is the concept from Web Squirrel [1] of agents. An agent is essentially a collection determined dynamically by a query. Membership in a Web Squirrel agent list is "by reference" — besides its listing in the agent list, a member has a "real" location and clicking on it in the agent list takes you to that location. Tinderbox [2] also has agents.

The kind of generic collections present in Web Squirrel or Tinderbox as agent lists are not "true" generic spatial collections, in that they don't have a full set of spatial attributes in their context in an agent list. The only attribute that is determined by query is collection membership. (Tinderbox allows the user to open a spatial view of an agent list, but members of the list cannot be manipulated spatially in the same way as a "manual" collection.) A fuller form of generic spatiality would allow for all of the spatial attributes that can be manipulated interactively to be determined by query criteria. Currently no spatial hypertext systems support this form of generic spatiality.

It may be argued that there is a total clash of principle between this concept of spatiality by query and the usual approach to spatial hypertext as the vehicle for implicit and emergent structure. Part of the motivation for spatial hypertext is the unwillingness of users to commit "in advance" to structure [8]. Here we are asking the user not simply to commit to structure, but to actually commit to a rule for how the structure is constituted. Be that as it may, the kind of queries found in agents such as those in Web Squirrel may be described as secondary in the sense that membership in the universe available to the query is already determined by a "spatial decision" of the kind familiar in spatial hypertext. E.g. a query might match the rule: "adjacent to an object whose name contains the string 'formula'". The question of what constitutes adjacency would presumably be determined by a spatial parser, which in turn is acting on the "raw data" originally provided by the familiar kind of spatial interactive placement. The exact relationship between implicit spatial structure and explicit query rules remains to be worked out; surely it seems wrong to "prohibit" agent rules from spatial hyper-

text due to an "ideology" of implicit structure. It will be up to users to determine how this gets worked out.[1]

EMERGENCE VS. CONDITIONALITY

Support for emergent structure has long been an important motivation for spatial hypertext. A natural component of the concept of emergence is completion: a structure may be incomplete, but is "growing". What happens if completion never occurs? Consider a collection containing summary thoughts for what is meant to be a section of a paper. The collection currently has two members. As the document author, you feel that two is too small a number for the finished collection, but aren't worried because you expect more members of the collection to materialize. However: if these members don't materialize, the members already in the collection need to be "reassigned". Perhaps if the emerging structure "fails", its current members should be assigned to the parent collection and their current collection deleted. Thus we have members of a collection where the membership is conditional on completion of the collection, and a specific behavior in mind if the collection can't be completed: the members are moved to the parent collection and the collection is deleted. Unfortunately, existing spatial hypertext systems have no way to indicate either this conditionality or the kind of behavior that should occur if the condition is not met. In the case we are discussing, as document author you must (1) decide that a collection has failed; (2) remember that you intended to move the members to the parent upon failure; (3) execute the failure behavior by hand.

By contrast, a spatial hypertext system that supported conditionality would allow you to have an attribute of a collection — call it "accepted", set the initial value to yes, and include in the collection a rule that if the accepted attribute changes to no, the members of the collection should be transferred to the parent and the collection deleted. Membership in the collection is conditioned on the collection being accepted. The entire hypertext could have a global constraint that when the entire hypertext has the attribute "closed", a collection must have at least 3 members in order to be accepted. Closing the hypertext would clean it of too-meager collections all at once. Note that for this to work correctly, it is crucial that the constraint requiring a collection to have at least 3 members be executed

depth first. (A collection with only 2 members may acquire more if one of those 2 is a collection which is "failed" under this scenario.)

CONDITIONALITY AND TIME

Many of the issues already discussed have to do with time. Given that VKB makes explicit reference to time as an "operator dimension", some explicit discussion of time vs. conditionality is in order. VKB allows the timeline to be set to the point at which an object changed. The assumption is that if an object changed, it was the user that changed it — presumably interactively. However if an object or collection can be subject to rules, it may be changed as the result of an interactive change to some other object. VKB records each change as an event. As rules cascade changes through other objects, do we have a single event or multiple events? Note that ordering becomes an issue here. Where all changes occur interactively, there is no ambiguity to ordering. However if an interactive change to one object results in changes to other objects through the application of rules, the order of application of the rules affects the time sequence of changes to other objects.

There are other issues related to time which are likely to be subject to severe aesthetic differences among possible authors. Suppose some of the objects that can appear in the same space are multimedia player objects of some kind — i.e. they have their own timelines, with activities programmed based on time. How should these timelines be synchronized? What kind of event handling framework is required to account for what should happen when objects "appear" or "disappear" as the result of conditional behavior? Time behavior could also be used as an "associating" attribute in a spatial hypertext, in the same way that such attributes as color and alignment are now. Consider a "scene" in a spatial hypertext in which several different objects are exhibiting time-based behavior. Some are in sync and some are not. The eye will clearly associate those that are in sync, in the same way that the eye will associate objects of like color. This type of association is likely to be extremely conditional, in that it may be related to when various objects were "set playing". Like color, the association by the eye of synchronized objects can occur at a distance. Far more subtle effects than synchronization are possible; the rhythms of time-based objects may overlay in slowly shifting patterns, similar to the rhythms in the music of composer Steve Reich.

Such "time structures" can be implicit, ambiguous, and emergent — exactly the kind of concepts that have been central to spatial hypertext since its inception. These effects can produce a kind of "layered time".

Of course some of these effects may be unintentional: objects may be synchronized simply by happenstance rather than an overt decision of the document author. Another complicating factor is that events may occur with a timing determined by the reader. If a spatial hypertext contains time-based multimedia objects, a great deal about the timing of when players are started and stopped will be in the hands of the reader.

CONCLUSION

Conditionality may be introduced into spatial hypertext by a variety of methods, from layering to explicit rules. As conditions change, objects may appear or disappear or be subject to other forms of behavior; this introduces new aspects of time to spatial hypertext which can interact with spatial aspects in complex ways. There can be odd effects in which attention attributes seem to be working at odds with structural attributes. Clearly, there is much room for further investigation here as more spatial hypertext systems acquire aspects of conditionality.

[1] It is interesting to note in this context that the explicit structuring of links has reemerged in more recent versions of VKB.

USER INTERFACE BEHAVIORS FOR SPATIALLY OVERLAID IMPLICIT STRUCTURES

INTRODUCTION

Spatial hypertext systems, such as VIKI [4], CAOS [5], and VKB [8] provide facilities for implicit or emergent structure by supporting spatial proximity as a form of association. The most intimate form of proximity is an overlay, in which items are placed on top of one another. Nevertheless, existing spatial hypertext systems provide only very weak support for overlays, since they assume that some part of each element in a spatial aggregate is always visible where it can be selected by clicking with the mouse. The problem, in a nutshell, is that when elements are overlaid an element can become completely occluded. Some explicit form of user interface behavior is necessary so that such elements can be discovered and accessed. This paper presents several such candidate behaviors. Most of these behaviors have in fact not been implemented; they are being presented here in the hope that spatial hypertext system implementers will find their ideas fruitful for incorporating into future systems.

In the examples that follow, figures illustrating user interface behaviors all refer to the same "baseline" spatial overlay, as shown in Figure 1. Figure 1 shows three VKB-like collections overlaid; the titles of two of them ("Anti-Conjunctive Drift" and "Sub-screening" are visible, while a tiny sliver from the third collection is just barely visible behind the other two. Within each collection are individual objects which are also heavily overlaid, so in fact there are 3 collections and 7 individual objects represented. Figure 1 shows only a single individual object readably; it is exactly how to discover and display the hidden objects that is the subject of this paper. Figure 2 shows the objects spread apart so that each one is visible. (Spreading is an important specific behavior, which will be discussed below.)

CURRENT BEHAVIORS

INDIVIDUAL SELECTION

All current spatial hypertext systems allow an object to be selected by clicking on some visible portion of it. Clearly this leaves completely unsolved the problem of selecting an object which is completely occluded. VKB allows selection of an occluded object by means of an object hierarchy view which is separate from the spatial view. This paper addresses methods by which occluded objects may be selected within a spatial view.

TRANSPARENCY

VKB allows the user to make objects transparent. Figure 3 shows the result of making all objects in Figure 1 transparent. As is readily seen, the result may be artistically interesting from a visual point of view, but does not contribute to the problem of navigating to and selecting specific objects.

DISCOVERY BEHAVIORS

Discovery behaviors alter the appearance of the visual workspace so that the location of hidden objects may be revealed; typically these behaviors do not show the entire contents of the objects revealed.

THE FRAME X-RAY

When the Frame X-ray behavior is invoked, a section of the visual workspace is selected and all objects in that section are rendered transparently, with borders intact but no "content". Figure 4 shows the results of the Frame X-ray applied to Figure 1. While the objects in Figure 3 are almost indiscernible one from another, it is surprisingly easy to pick out individual objects in Figure 4. Note in particular that the hindmost collection is quite easy to spot. The idea of the Frame X-ray display is that object borders are available for selection even on objects that are many layers down.

While the Frame X-ray concept is very simple to understand and should be fairly easy to implement at the application level, it does have some clear drawbacks. (1) While the existence of objects is made clear, they are not differentiated one from another by content. (Though border color distinctions would be available.) (2) The behavior it introduces is likely to be highly modal. An object whose outline is revealed in a Frame X-ray display must be selectable. I.e. interactivity of

the visual workspace must be preserved with some degree of orthogonality with respect to Frame X-ray display vs. normal display. This complicates the user interface; designers may have strong opinions about modality. (3) The Frame X-ray display itself is confusing if objects are highly aligned geometrically.

THE CORE-DRILL

The core-drill is a moveable geometric region of the screen — presumably a rectangle — which has popped up next to it a display of names of all the objects it intersects. This is illustrated in Figure 5. When an object from the menu is selected, it is brought to the top; if that object is part of a collection then presumably the whole collection is brought to the top. Note the menu in Figure 5 is shown as "smart" about collections: objects inside a collection are shown indented under the collection name.

The core-drill allows a very effective form of navigation of the visual space, in spite of a great density of overlays; if the associated pop-up menu has scroll bars, then there is no limit to the amount of overlaying that can be navigated. However, there is one obvious drawback to the core-drill: to be effective, all objects must have names. This is a very major issue. Spatial hypertext was created in the first place to avoid the problem of premature commitment to structure [3]; premature commitment to naming is not all that different from premature commitment to structure. Figure 5 shows some names that were "manually" created. Of course an agent could create names for objects dynamically — which might work well if all objects are text objects — but there remains the problem of how names would be created for non-text objects, such as graphic images or multimedia objects.

The core-drill also has modal issues; presumably there is an interface behavior or keyboard shortcut that causes it to appear and disappear, and modality of dragging around or resizing the core-drill vs. "normal" behavior of the spatial workspace could be tricky.

THE GLASS-BOTTOM BOAT

The Glass-Bottom Boat is a special window which "slices through" the layering of objects on the visual workspace to reveal whatever may be present several layers underneath the surface. This window has explicit gadgets to raise it or lower it one layer, and an indication / setting to show how many layers down from the surface of the workspace it is. The Glass-Bottom Boat is quite attractive from a modality point of view: presumably it acts just like any other window, and within the sub-layer it reveals in its contents, events are passed through exactly as if that layer was on top. I.e. user interface behavior inside the Glass-Bottom Boat window is identical to behavior were the same contents to be part of the top layer. In this sense the Glass-Bottom Boat is not modal, and uses no interface resources that might be desired for some other purpose. Alas, this very power reveals the main weakness of this concept: in order to be implemented properly it might have to be located in the native operating system windowing system itself. Commercially viable windowing systems are not normally accessible to such interventions. For windowing systems implemented at the application level (e.g. Smalltalk) it might be feasible, however. The Glass-Bottom Boat is illustrated in Figure 6.

In spite of the simplicity of this concept, there are likely to be some surprises if it is actually implemented. Consider a Glass-Bottom Boat window that is fairly large, and which is several layers down, as shown in Figure 6. Suppose an object is selected from inside the Glass-Bottom Boat. To what layer should the Glass Bottom Boat window move? If (as discussed above) the interface inside the Glass-Bottom Boat works "identically" to the normal interface, the selected object moves to "the top". What is the top? Is it the top with respect to the current position of the Glass-Bottom Boat, or the top of the entire windowing system? If the object moves to the top of the entire windowing system, then we have the paradox that it might disappear from the Glass-Bottom Boat! Should the layer of the Glass-Bottom Boat then follow the selected object? Should the Glass-Bottom Boat automatically close? User preferences may need to be consulted to answer such questions.

EXPOSURE BEHAVIORS

Whereas the discovery behaviors discussed above tend to simply identify where a possibly occluded element in an overlay may be found, exposure behav-

iors tend to expose the entirety of an element. (Actually the Glass-Bottom Boat is a hybrid that could be classified either as an exposure behavior or a discovery behavior.)

THE FRAME STACK

When the mouse approaches an overlay having this behavior, a stack of frames is displayed, one for each element of the overlay. These frames serve as "On MouseOver" hot-spots to allow navigation of each member of the overlay; note that the frame stack explicitly embeds the peer structure of the overlay. The frame stack is illustrated in Figure 7. The frame stack concept has been implemented in specific literary works [6], though no authoring environments make it easy to implement as an "off-the-shelf" behavior. The frame stack concept works well with nested collections, since it allows a frame stack to be opened inside an outer frame stack. One drawback to the frame stack concept is that it is quite expensive in terms of screen real estate. Assuming all frames in one stack are the same size, this must be large enough to accommodate the largest element — with sufficient room for a bit of margin. Note that as rendered here, if a frame stack were to be opened on the contents of the collection shown on top in Figure 7, there would not be enough room to fit all the frames inside the current collection boundary.

THE SPREAD

This behavior was explored by Mander et al [2]; the spread simply spreads out members of a pile so that all of them are visible. Figure 2 shows a "recursive spread" of the example pile used throughout this paper. Spreading uses a technique which may be called co-presentation, which is discussed in detail in [7]. While the simplicity of spreading is appealing, there are a number of issues with this behavior. It appears that for the foreseeable future, screen real estate will always be in short supply; what happens when a pile needs to be spread out over more real estate than is available? In this case we will still have members of the pile potentially occluding one another, which means some other behavior will be needed. When piles are nested, presumably each collection would have to occupy a disjoint space (as in Figure 2) with spreading inside that space. This explodes the space requirements significantly.

THE VIEWING CONE

Mander et al presented another idea for navigating piles, which they termed a viewing cone. Their implementation of piles showed a collapsed pile as a 3-dimensional stack so that an edge from every member of the pile is visible. The viewing cone expands a particular member of the pile to a thumbnail or larger view. Again, it is not clear how this would work with nested piles, and also has the problem that at least some part of each member of the pile must be visible; the viewing cone cannot locate an object if it is completely occluded.

GEOMETRIC RECTIFICATION

There are many ways this behavior can be implemented; it is really a special form of spread, except that parts of objects may still be overlaid. At least one form of this behavior is widely available: many multi-window GUI applications have a "cascade" option for viewing windows, that arrays the windows so that their title bars are all visible as a descending "slant". This behavior has many of the issues of spreading.

GENERAL BEHAVIOR ISSUES

There are several general issues that pertain to all of these behaviors.

SCREEN GEOMETRY REQUIREMENTS

In a system like VKB, a collection is assumed to have a particular size, which is a property of the collection; it can also be zoomed to occupy the full available screen area. It is clear that several of the behaviors discussed above may have additional conditional requirements for screen real estate. Spreading requires enough real estate for all items to become exposed; the frame stack requires enough real estate to show a frame for each element of a pile, etc. How does an object such as a pile communicate to the user interface system how much real estate it requires for some behavior which will only occur when the user initiates some particular action? How does the user interface system communicate to an object how much screen real estate is available?

PLUGGABILITY

Ideally, user interface behaviors of the kind discussed here should be pluggable; as new behaviors are invented it would be wonderful if they can be implemented in existing systems without having to completely reimplement a system. Of course this requires a framework, which is a significant effort. The Model-View-Controller Paradigm [1] represents a classical effort to implement pluggable user interface behaviors; it is complex and in spite of frequent revivals often not used.

Spatial hypertext systems may have special issues with regard to pluggable behaviors. The whole rationale of spatial hypertext is based on supporting emergent and ambiguous structure. An object may be placed "near" another object to indicate an ambiguous or not-yet-defined relationship to that object; this association might not be represented by any persistent structure, but might be computed on the fly, based on criteria that can change depending on what the user does. To the extent that user interface behaviors have geometry requirements, such requirements may have interactions with such on-the-fly computations. How do these communicate? For example: computation of available real estate may be affected by what counts as "near" in spatial parsing algorithms. If a spread is supposed to avoid overlaying a nearby collection, the available real estate for the spread is constrained not only by the actual real estate occupied by the nearby collection, but by a surrounding area which the spatial parser would "count" as part of the collection. Thus a spatial parser would have to respond not only to requests to indicate what objects are "in" an aggregation, but also what areas include objects that would be considered "in" an aggregation just by being present in that area.

BEHAVIOR NEGOTIATION

A common theme of user interface behaviors is that space requirements can vary depending on the state of an object. An object which is under active investigation may be "expanded", and when not receiving attention may be "collapsed". Current spatial hypertext systems do not support negotiation of objects for screen real estate. For instance, in VKB it is assumed that when a collection needs attention, it can be expanded to fill the whole screen — eliminating from

view “sibling” collections. A system of negotiation would allow an object that needs more screen real estate to obtain it by a variety of means. If expanding the collection to the full bounding box of all objects in that collection does not collide with the space of any other object, it could be expanded to just that amount of space, with no impact on other objects. If not enough real estate is available, nearby objects could be “asked” to collapse. Finally, if insufficient real estate is still not available, objects could be overlaid. It goes without saying that the design of a protocol whereby such negotiation would occur is a significant undertaking.

The concept of spatial hypertext includes the idea that certain objects need to be seen “together” — if at all possible. This calls for some subtlety in the design of the user interface when these same objects (alas) compete for user interface resources. The requirement of objects for “coattention” has unexplored consequences for user interfaces.

NAVIGATING NOWHERE/HYPERTEXT INFRAWHERE

MUST THE LEXIA BE LINEAR?

It is a truism of hypertext rhetoric that navigation is the user's choice. If we truly believe this, then among these choices we must place the null choice: the choice to go nowhere. This presentation begins with the rhetorical question: "When you go nowhere, where are you?" The conventional term in literary hypertext for "where you are" in a hypertext when you haven't gone anywhere is lexia (Hypertext fn). Thus we begin with a discussion of the structure of the lexia. It is customary to avoid this question: the lexia is typically considered ordinary linear writing, with the "real" issues of hypertext relating to how lexia are organized; the lexia itself is not really considered as a hypertext. (E.g.: "[The within-component] layer is purposefully not elaborated within the Dexter model" (Halasz and Schwartz 32); "The local stability of the lexia arouse expectations of coherence and internal consistency" (Moulthrop, "The Shadow of the Informand.") However, non-linear writing, particularly in poetry, has a rich and varied history; poets who are comfortable with extending that history have reason not to accept an inevitable linearity of the lexia.

Intergrams illustrates an example of a non-linear, non-link structuring method, the simultaneity: the literal layering on top of one another of language elements. (See Figure 1.) The hypertext link does not really express this structure, particularly as it is actually realized by existing user interfaces. A much closer match is provided by the concept of an Aquanet relation (Marshall et al. 265) [1]. A simultaneity may be thought of as a relation in which all slots are unnamed, equivalent, accessible, and located in approximately the same graphical position on the computer's display. In Intergrams, simultaneities are implemented through the use of "tactile", no-click buttons in which there is screen behavior from simple mouse-cursor movement through hot-spots on the screen. Such devices may have uses beyond the artistic. The concept of simultaneity is the direct

analog for writing of juxtaposition concepts which have a rich and varied history throughout the twentieth century in visual arts and music, and is clearly related to simultaneous-voice work (both live and on magnetic tape) in poetry. (See e.g. [Ros75a].)

Figure 1: Screen dump from Intergram 9. The top rectangle shows a simultaneity which is "closed" — all phrases are visible simultaneously — while the bottom rectangle shows a simultaneity which has been opened to show only one phrase. The rectangle on the right is a button serving as a link to a relation among three simultaneities; the button at the lower right is linked to the parent screen. This screen shows a ternary syntax relation with the "verb slot" (to use Aquanet terminology) toward the right.

Another example of how the lexia may be structured nonlinearly is polylinearity: the stringing of word skeins in a graphical space where normal print conventions establish no clear ordering among the skeins. The simultaneities in [Ros94] include both polylinear text and "single-card" relational syntax diagrams. (See Figure 2).

Figure 2: Polylinear text from Diffractions through: Thirst weep ransack (frailty) veer tide elegy. This is only one plane in a simultaneity consisting of 2 other polylinear screens and 3 relational syntax diagrams.

Is "the lexia" a single plane in a simultaneity, or the totality of all planes stacked together? In a work like Intergrams, one could easily argue there are "conventional" lexia, that a single plane in a simultaneity acts as the lexia. The same could be said of Diffractions through, though in that case the lexia is fractured by polylinearity and sub-diagrams. Is "the lexia" in a polylinear case a single word skein, or the whole net? This question is a bit more troublesome. One could be

comfortable arguing tha "the lexia" is whatever you see on the screen without moving the mouse or touching the keyboard. On the other hand, given a representation of a hypertext as a network of nodes, the traditional concept of lexia is: one of the nodes. Should we insist that lexia be node-focused in the case of a typical hypertext but insist on accepting all skeins simultaneously in the polylinear case? This seems arbitrary.

CONJUNCTIVE VS. DISJUNCTIVE STRUCTURE

Notwithstanding the celebrated remark of J. Yellowlees Douglas [Dou91]:

They are all laid out before us: the genuine post-modern text rejecting the objective paradigm of reality as the great "either/or" and embracing, instead, the "and/and/and."

in fact, the typical hypertext link may be described as a disjunctive link: if lexia X has links A, B, C, D, the user may choose A or B or C or D (or to go nowhere, of course!). Almost the entirety of hypertext rhetoric surrounds what may be called "the confrontation with or" — how to assist the reader in coping with the volume and structure of choice. A simultaneity may be disjunctive or conjunctive: the whole of a simultaneity with planes A, B, C, D may be A and B and C and D. One might envision hypertext links as being conjunctive also. Surely in many cases Aquanet relations must be described as conjunctive, not disjunctive. Consider, for instance, an Aquanet schema for diagramming sentences in which a Sentence object has slots for Noun Phrase and Verb Phrase. May we describe the verb phrase slot as optional?

While the literature concerning the rhetoric of the hypertext link has become amazingly voluminous considering how few literary hypertexts there still are, the rhetoric of the conjunctive hypertext relation is nearly non-existent. Discussions of hypertext consistently use travel vocabulary: links are "followed"; [Lan89] speaks about "arrivals" and "departures" — all terms that are quite cogent when applied to disjunctive links, but which may have limited relevance when applied to conjunctive relations. Perhaps we should speak of gathering a conjunctive relation, rather than following, as in the case of a disjunctive link. This has

major implications for discussion of the lexia. "Following" a link means leaving it for another lexia; "gathering" a relation means bringing things to a central place: whereas the disjunctive link is associated with travel, the conjunctive relation is associated with locus, with an inherently structured lexia. Stuart Moulthrop, in [Mou92b], concurring with [Bol91], seems to apprehend the problem, but has proposed the wrong solution: a kind of Wittgenstein duck-rabbit flip in functionality between node and link. While this is a fascinating and compelling metaphor, the structural inadequacy of having nothing but nodes and disjunctive links is not to be solved by having nodes and links philosophically trade places but by the much more obvious expedient of providing more explicit structure. (But Moulthrop and I come out the same place in the end, see below.)

That the concept of relation is closely tied to spatialized text has been studied in detail by Marshall and her colleagues ([Mar92], [Mar93]). She relates that Aquanet users frequently constructed piles where it had been anticipated they would construct relations.[2] Perhaps systems such as Aquanet should formalize piles as legitimate objects. Certainly the pile is a worthy artistic device in its own right. Piles are in widespread use in the visual arts, and e.g. [Ros73] and [Ros75b] used word piles as a formal device, in the latter case including word piles into higher-level relational structures.

The proper user interface behavior in the face of conjunctive structure is an interesting challenge. It is likely that writers will want different word object behavior from conjunctive structures than from disjunctive structures. Note this has obvious ramifications for formal theories of hypertext. The assumption is explicit in [Gr¿ 94] that the Dexter concept of composite adequately formalizes Aquanet relations, and there is an implicit inference in [Mar91] that hypertext links are special cases of Aquanet relations. Composites are arguably conjunctive, but it is not clear that the conjunctive relation has the same object behavior requirements as the composite (which may have no behavior requirements at all), and certainly disjunctive links and conjunctive relations are vastly different kinds of objects, and deserve to be recognized as such both in software object models and formal theories. Indeed, as a writer, I would make an urgent plea to software developers to include as many object types as possible in your models, and abso-

lutely to provide some form of extensibility so that writers can program their own way to workarounds when off-the-shelf object behaviors are not adequate.[3]

The concept of contour, [Ber92], which I would describe as a supra-lexial projection by the reader of a geography underneath pathways through the orcloud, is perhaps intermediate between the purely disjunctive and purely conjunctive. Can we describe the contour as the attempt to resolve disjunctive experience into conjunctive resonance?

FROM INTENSITY OF STRUCTURE TO HYPERTEXT INFRAWHERE

The author's current work in progress (as yet untitled) carries the concept of simultaneity still further in the idea of a nested simultaneity. In some cases this work carries the simultaneity inside the sentence. Hypertext is carried into the fine structure of language. Where is "the lexia" now? Is there really a concept of lexia when we are inside the sentence?

A hypertext may be thought of as a kind of virtual diagram, with software for navigating the diagram. If the diagram is small enough it may be presented in a single graphical space, without the aid of software. The author's Diagram Poems, e.g. [Ros79], [Ros84], are examples of such works. These present an explicitly relational syntax notation, still used in Intergrams and Diffractions through. The structural atoms in the Diagram Poems are small clusters of words; the relational (i.e. hypertext-on-paper) structure is the sentence structure.[4] What shall we say is "the lexia" here? In the Diagram Poems, the diagram notation carries syntax itself. Executed on a larger scale, this concept leads to the use of hypertext to carry the very infrastructure of language. Such works would have hypertext infrawhere: a structural underneath so fine and so pervasive, a lexia so completely fragmented, that the concept of lexia ceases to have any meaning: a completely dematerialized lexia, as in [Mou92b] after all.

In [Mou92b] Stuart Moulthrop asks: "Why does the hypertext research community publish its work in print?" At the risk of seeming glib, the answer is obvious: because hypertext is not our native tongue. Many will surely balk at

the idea that this needn't be so, that there can exist a natural language in which hypertext carries the very structure of syntax itself: hypertext not as a medium of organizing thoughts, but as a medium of thought. Perhaps in the end this will turn out to be unachievable, but as a focus for poetic experimentation it provides this author with a sustaining vision.

[1]Not having had the pleasure of being able to personally use Aquanet (alas!), I am relying on published descriptions of this software.

[2]It may be thought that 'pile' and 'simultaneity' as presented here are the same concept. They are very closely related, though apparently the piles described in [Mar92] were spatially aggregated units with no actual word object behavior, whereas the simultaneities in [Ros93] have explicit well-defined behavior.

[3]Intergrams is implemented in HyperCard; the simultaneity abstraction is implemented painfully and "by hand" — there is no such abstration "built in" to HyperCard. While I would have many criticisms of HyperCard — it is for instance very poorly orthogonalized (a button is a "pluggable" object but cannot have storage containers; a card can have storage containers but is not pluggable; thus pluggability and ability to have storage containers are not properly orthogonal) — still it is a tribute to HyperCard's flexibility that such an abstraction as simultaneity could be implemented in it with only a modest amount of effort.

[4]The notation in the Diagram Poems allows the use of such devices as internal links (links between a component and a larger whole in which it participates) and feedback loops — constructs which the eye can resolve into a gestalt but which may present software and formal theories with some difficulties.