Philadelphia Fullerine: A Case Study in Three-Dimensional Hypermedia

J. Nathan Matias
Elizabethtown College
One Alpha Drive
Elizabethtown, PA 17022 USA
1-717-653-4669
Email: matiasj@etown.edu

ABSTRACT
Philadelphia Fullerine, a geodesic hypermedia sculpture designed by the author, is about ethnic and lower class life in mid-19th century Philadelphia. Each of the 60 faces presents primary image material and a short audio documentary. Adjacent faces are linked conceptually. This geodesic sphere has full rotational freedom. Viewers are encouraged to begin anywhere and follow any path of adjacency. This paper examines the underlying theory, design methods, and structure of the sculpture as a case study in the applications and challenges of creating, storing, and navigating three-dimensional hyperstructures with spatial hypertext software and GZigZag.

Categories and Subject Descriptors
H.5.4 [**Hypertext/Hypermedia**]: Architecture, Navigation, Theory, User Issues.

General Terms
Design, Experimentation, Human Factors, Theory.

Keywords
Authoring, creative nonfiction, directional links, Gzz, hypermedia topology, implicit structure, information triage, sculpture, sculptural hypertext, spatial hypertext, transclusion, Tinderbox, ZigZag, zzstructure

1. INTRODUCTION
Hypermedia projects often attempt to represent virtual information through physical metaphors mediated by computer hardware. The hypermedia sculpture *Philadelphia Fullerine* turns physical metaphor into a physical reality. This hypermedia is a 60-faced geodesic sculpture, paired with audio.

This work has provided a case study in the use of spatial hypertext and sculptural hypertext methods to author a research-based creative nonfiction. The hypermedia also provides a case study in the encoding and navigation of three-dimensional, geometrically-based hyperstructures. By analyzing these tasks, I hope to raise questions and encourage more discussion of the practical implications of certain user interface design and hyperstructure concepts.

HT'05, September 6–9, 2005, Salzburg, Austria.

2. DESCRIPTION
Philadelphia Fullerine is a creative nonfiction hypermedia consisting of 60 triangular images and 60 segments of audio documentary. The hypermedia presents images, people, stories, events, and information about mid-19th century ethnic life in the city of Philadelphia. Each triangular node links to three adjacent triangles, forming a geodesic structure.

The images are placed on a 60-face, 32-vertex geodesic sphere, mounted inside two rotating steel rings which spin on a pedestal. Three axles are arranged on the concentric rings and base at right angles to allow the sphere full rotational freedom. The sphere's structure consists of a framework of steel rods and connectors. The structure is wrapped in coarse jute string. Handles screwed to each of the 12 pentagonal joints provide a basic way for users to rotate the sphere. The balsa faces of the sculpture, which vary in shape, are suspended by hemp cord between the triangular gaps formed by the structure of the rods. These faces mostly depict images from the 19th century, though some recent photographs and digital artwork are also displayed. The wood and coarse textile materials, organized in a 20th century structure, symbolize the mid-19th century tension and transition between handmade culture and manufacturing society. Resting on its base, the sculpture is 1.4 meters tall. The sphere and concentric rings are approximately 67 centimeters in diameter.

The sculpture requires the use of a supplied CD player. Each image is paired with an audio segment which varies from one to five minutes long. To help users find the corresponding audio, triangles are labeled with a track number. The total audio content amounts to two hours and twenty minutes on three compact discs. Although inefficient for the user, compact-disc technology was chosen as the most cost-effective audio delivery method.

Adjacent triangles relate in chronology, theme, detail, or characters. Since it is spherical, the sculpture suggests no beginning or end, Users may begin anywhere, end anywhere, and follow any desired path through the visual and audio material. Rarely do adjacent triangles share adjacent numbers. The numbering has been specifically designed to discourage a sequential, numerical encounter with the material.

Figure 1. Philadelphia Fullerine

2.1 Goals and Artistic Vision

Life is. We attempt to catalogue it, lock it up in binders, link it together, or otherwise bend and extract it into meaningful chunks and structures. History is one class of such thought, formed from our collection of stories, events, causes, prejudices, people, coincidences, suspicions, observations, analyses, correlations, preferences, assumptions, and best efforts to understand.

Philadelphia Fullerine does not attempt to escape the difficulties of assembling irregular information into a coherent whole. Instead, it displays the inadequacies of objective historical thought while maintaining an honest attempt to understand the complex issues of life for ethnic groups in mid-19th century Philadelphia. In a time of great reform and social optimism, great violence, prejudice, and confusion remained starkly prevalent. On one hand, advancements in education, industry, and infrastructure increased quality of life in the long term. Yet these improvements caused great economic troubles and encouraged greater injustice for decades.

Rather than presenting a single, fully-cohesive understanding, this sculpture attempts to educate while subverting a neatly-conclusive view of history. First, this narrative's presentation as a hypertext suggests the difficulties with truly linear views of history[3]. Secondly, its physical structure reminds users of the complex interactions of life's whole experience. The geodesic structure further subverts its own seeming authority, because a neat, highly geometric understanding of history is at least as unlikely as a linear structure. This infinitely-recursive structure further subverts a linear, authoritative view of history by giving viewers no clear beginning or ending point. Finally, by presenting more audio information than users are likely to experience in one encounter with the sculpture, the project emphasizes the limitations of our ability to know.

The inconclusive nature of the structure is not intended to discourage viewers, but rather to free them. The spherical shape frees viewers to begin with whatever image catches their interest. The promise of relationships between adjacent nodes encourages them to explore and discover. They may conclude where they wish. The sculpture makes fewer demands upon the viewers than linear writings. Viewers need not feel guilty for not finishing the narrative or for not following arguments to a conclusion. Among this subversion of a linear, conclusive view of history, the spherical shape and hypertext connections suggest a sense of wholeness. Viewers are encouraged to think about history as more than a single event or topic, but rather as a combined experience which includes events and details from all aspects of life.

As a hypermedia, the sculpture encourages viewers to critically consider the information presented in the sculpture. Links are based on adjacency. Links do not describe the nature of the connection, nor are they directional. The author merely suggests the presence of connections through adjacency. The meanings of connections must be determined by the viewer. Some adjacent nodes even contain contradictory information and opinions. Rather than steer viewers toward a interpretive framework of historical events, nodes with clear, well-reasoned conclusions are sometimes placed adjacent to nodes that present information that may seem to undermine those conclusions. Since no particular path is guaranteed, different viewers may gain completely unique understandings of the contents and meaning of the sculpture.

This project was intended to be part of my research for a larger, traditional historical narrative text on ethnic life in mid-19th century Philadelphia. I hoped that the act of authoring a hypermedia would force me to study a wide range material outside the events of my planned narrative. Research for the hypermedia informed my linear writing with insight into the texture of life for people during the my target time period.

Since I intended to create two projects from the same research, I needed a good way to store and organize my findings. The project provided me with a further opportunity to refine the hypertext research and writing process I have developed in Tinderbox[12]. By conducting all research, planning, and writing in Tinderbox, I surmised that my research efforts would become more efficient and coordinated. Furthermore, I intended to use this deeply-interlinked research to experiment with link culling techniques[1] by forming the geodesic hypertext through the removal of links.

2.2 Geometry

Philadelphia Fullerine is not the first multidimensional geometric hypermedia. Rather, structure seems to be a central issue to many works of hypertext. Many hypermedia works center around the influence of structure to aid or disorient the reader[5]. As demonstrated in Parunak's analysis of hypermedia topologies, geometric organizations of information have long been suggested to aid a user's ability to find specific information within hypertexts[13]. In more complex systems like ZigZag, the path strategies outlined by Parunak are used to navigate information with non-euclidian structure[9].

[1] To avoid confusion, the term "culled" will be used to refer to the methods of "sculptural hypertext." For further explanation, see section 3.1.

"Letter to Linus," a hypertext poem by William Gillespie, is one example of the hypercube geometry suggested by Parunak. In this cube-shaped poem, each side represents a different canto.. One may begin the poem from any side. After reading the canto, the reader is shown a representation of the cube, with its sides folded out so the titles printed on the adjacent sides are clearly visible. When the user clicks on the title of an adjacent side, the cube is rotated so the newly-selected side now faces the reader, and the corresponding canto is displayed[6]. A representation of the newly-rotated cube is presented for navigation.

Many geometric hyperstructures, including the hypercube and hypertorus described by Parunak, do follow a matrix-like layout, with links along the x and y axes. A hypercube structure like that of "Letter to Linus" would seem to provide a simpler geometric hypertext than a geodesic sphere. Indeed, I originally intended to use a hypercube structure. Each side of the cube was to contain nine square nodes. However, as this structures scale in nodes, it becomes increasingly complex. Each node would link to four adjacent nodes but would also be part of six major groupings which correspond to the sides of the cube. This would add eight implicit links to each node. Implicit links would have vastly increased the number of connections required for a successful work. Influenced by these groupings, users would also be less likely to consider any node as a potential beginning or end.

I chose a Type II geodesic structure of 60 faces instead of a cube. Implicit, spatial groupings can be derived: the sphere can also be thought of as twelve pentagons of five triangles. However, these groupings are difficult to identify on the sphere. Irregular hexagons may also be seen in the sculpture, but they share triangles and are also difficult to distinguish.

Instead of using squares in an orthogonal structure, the geodesic sphere uses triangles, which require fewer links of adjacency per node than squares. Surprisingly, the geodesic structure requires a greater total number of adjacency links than a hypercube. A cube of 54 nodes only requires 108 links at four links per node, resulting in 2 unique links per node. The geodesic structure of 60 nodes requires 140 links at three links per node, resulting in 2.3 unique links per node. These ratios do not include the implicit links that would be suggested by the groupings on the side of a cube. The complexity of a cube is much greater than that of a sphere when these implicit groupings are considered.

By using triangles in an arrangement that discourages implicit, spatial grouping, the mental effort of assembling the hypertext became easier. Since each triangle links to only three other triangles, changes to a single node affected fewer nodes than similar changes in a cube.

Issues relating to electronic representation and user navigation of this geodesic structure are examined in section 4.3.2.

3. RESEARCH, DESIGN, AND CREATION

Spatial hypertext software was critical in the design of this structured hypermedia. It was also useful throughout the research and writing stages. By storing all information –sources, research, ideas, text, tasks— in hypertext software, the creation of the sculpture and its structure became much easier.

3.1 A Note on Terminology: Sculptural Hypertext or Hypertext Sculpture?

Philadelphia Fullerine is a physical work of mixed media sculpture. It was also partially designed using methods of "sculptural hypertext." Bernstein's term "sculptural hypertext" refers to hypertexts whose narrative structure has been shaped primarily by the removal of links. He compares the creation of such hypertext works to the act of carving stone sculpture from solid materials[1]. Most sculptural hypertexts do not exist as a physical work of sculpture artwork, but .as electronic documents. To avoid confusion, the term "culled" and related conjugations will be substituted for the term "sculptural" when used in the context of the link-removal hypertext authoring technique described by Bernstein[1].[2]

3.2 The Role of Hypertext Software

Spatial hypertext tools have been shown to aid users in managing and making decisions from large amounts of information[8]. I have also conducted and described a further case study showing the usefulness of researching large nonfiction writing projects with spatial hypertext[12]. With this in mind, I used Tinderbox, a spatial hypertext tool, to plan and carry out the research, writing, and construction of the sculpture.

In Tinderbox, users create notes, and "incorporate those notes incrementally and naturally into an organic, unplanned spatial hypertext." Custom data fields may be associated with these notes, which are the basic unit of data in Tinderbox. Notes may be the parent of other notes in a hierarchy. *Aliases* of a note allow a single note to take multiple places in outlines and maps. *Agents* search the document and aggregate groups of aliases according to user-defined criteria[2]. The data inside Tinderbox may be viewed in hierarchies, as spatial maps, as text with links, and several other ways. For example, the *Roadmap* view displays the contents in the center and shows inward and outgoing links in separate lists on the side of the window.

All potential sources were entered into Tinderbox. Sources were grouped by library. Each library's grouping was also a node, containing information about the library, its librarians, and its procedures. In Tinderbox, nodes may also be given custom data fields. I used these fields to specify common source information, such as author, publisher, due date, and a relevance rating. Excerpts and annotations were linked to these source nodes. Custom Tinderbox agents were used to search and aggregate sources on these fields, constructing lists of sources consulted, sources checked out, sources deemed useful, and rejected sources. When possible, the full text of sources was included in the Tinderbox hypertext.

To maximize the efficiency of research trips, I created task lists inside the Tinderbox space, linking tasks with libraries, museums, people, and bibliographic information about sources not yet obtained. I was thus able to stay focused and accomplish my research goals by storing and interlinking histories of past

[2] Sculptures may be made through many more techniques than the removal of material: the term "culled hypertext" might more clearly describe the linking practice described by Bernstein. The term is also more linguistically versatile

interactions, goals for research trips, listings of library holdings, maps, addresses, and other disparate, useful details.

Spatial interfaces allow the quick, meaningful inclusion of very unstructured data by means of implicit structure [8]. Thus, I was able to separate the initial entry of information from the task of refining explicit research-related links. Yet the research and planning stages occurred in parallel. As I found ideas for potential inclusion into the sphere, I placed them in a spatial map inside the Tinderbox file. I then linked them to all other potentially-related items in this map, using link types to describe the nature of the connection (see figure 2). As ideas for audio text occurred to me, I also wrote scripts in another part of the same document. These scripts linked back to the sources. During this time, I also worked to obtain digital proofs of potential images. These were catalogued as sources in the spatial hypertext document.

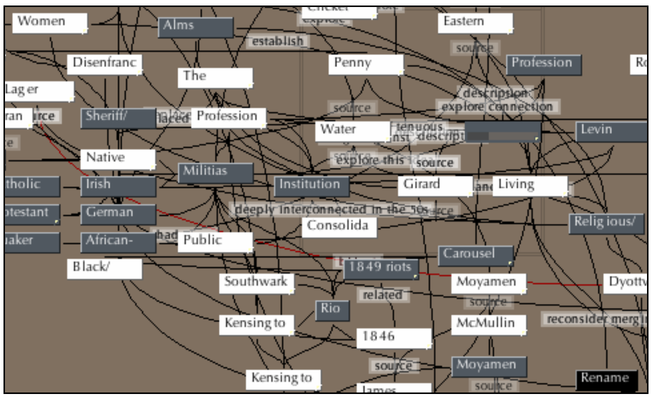

Figure 2. An excerpt of the spatial map of all potential links. White nodes were included in the final hypermedia.

The contents of the sphere were derived from the heavily interlinked collection of ideas and topics shown in figure 2. Decisions about node placement were determined by the structural requirements of the sphere, the availability of related images, and the limitations of my existing research. By using a set of initially-linked nodes, I culled from a good set of strong links rather than creating weak connections for structural reasons. Such weak connections exist, but the sculpture would be plagued by weaker connections if I had not used link culling techniques.

To avoid direct modification to my extensive set of possible connections while culling, I created aliases from my collection of linked ideas and placed them in a new container. The spatial view only shows links made to an individual instance a node. An alias of that initial node may be created and placed elsewhere, but link lines to the original node do not now also point to the new alias. Thus, if two nodes are linked in their original instance, a link line will not be drawn between their aliases. This feature gave me the power to experiment with a number of arrangements of the sphere without the risk of overly modifying my source collection of ideas or the need to create duplicate nodes.[3]

In Tinderbox, I was able to start with a highly-linked collection of information, culling links and nodes toward the final geodesic structure. The availability of numerous Tinderbox views of the same data was helpful during this culling process. As I culled and constructed the geodesic structure, a scale model of the sculpture also became very useful. While spatial hypertext features provided a way to encode, link, and develop the hypermedia, such two-dimensional views could not properly visualize the structure. Sticker labels on the model were used to synchronize the model and the contents of the Tinderbox map.

Sometimes, a careful look at the pre-existing links between nodes informed my decision about the placement of the next node. Other times, by viewing the progression of topics across the face of the scale model, I could strategize overall placement. I would then synchronize the Tinderbox space with the model (see figure 3).

Sometimes, pre-existing links were not sufficient. In those cases, I would split nodes or create new links. Each time I made a decision, I double-checked the availability of images and sources. This was easy, since all the data was within the same hypertext. The rich, interlinked dataset also simplified the task of compiling lists of image requests from archives and libraries.

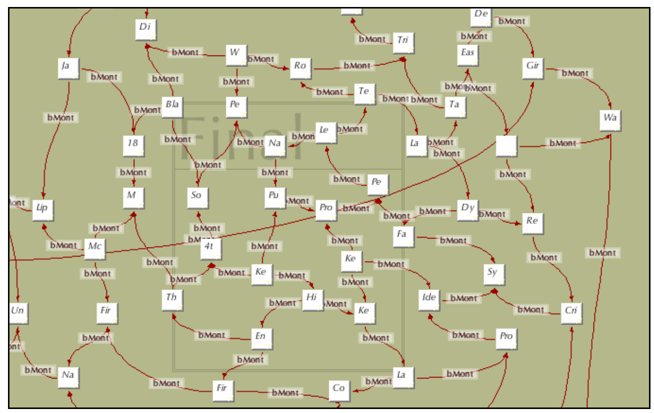

Figure 3. Part of the final, culled geodesic structure, consisting of aliases whose original instances are shown in figure 2.

The writing, recording, and editing phases of the adjoining audio documentary were also managed in Tinderbox. Sorted task lists and color-coded project-status views were indispensable for each of these stages. The power of hypertext software was in more than its ability to store and relate my research and creative data. It was able to encode, track, and interlink useful information about my efforts to work with the data I was storing in the same hypertext.

4. OBSERVATIONS AND ISSUES

Both the final sculpture and its process of construction provide insight and raise questions in a number of areas relating to hypertext. The following section is intended to contribute questions and inform further discussion rather than answer issues outside of the scope of this case study.

4.1 Aliases, Transclusion, and Links

Aliases are the equivalent of transclusion in spatial hypertext systems, even when nothing links to them. Marshall and Shipman have established that nodes in a spatial hypertext need not be linked explicitly to be linked. [7]. Their research shows that, for spatial hypertext, a two-dimensional area can be treated much like a document, since it may contain its both content and links, even when those links are implicit. Thus, by creating an alias in a spatial hypertext, one effectively transcludes an entire node.

[3] The relation of this feature to issues of transclusion and spatial software is considered in more detail in section 4.1

Tinderbox aliases allow a single node to exist in more than one location. They share the same title, text, and attributes with the original. In Tinderbox, which does not manage textual transclusions or version information, most changes to an alias also change the original node. However, individual aliases do have three unique properties: size, location, and links.

Location and size must be unique for each alias if they are to have spatial independence from their parent. Alias-specific links are also useful. In a multi-user system, a user with read-only access may wish to transclude(alias) something controlled by someone else. This user should still be permitted to interlink this transcluded item. Even in a single-user system like Tinderbox, giving aliases their own links becomes very useful. In *Philadelphia Fullerine,* this feature simplified the process of culling links from a heavily-mapped set of nodes. I was also able to use aliases to experiment with several different arrangements and cullings of the same data without needing to fully duplicate nodes.

First, I created a set of aliases from the collection of original nodes and links. Then, while viewing a map of the aliases and a list of links, I reproduced among the aliases all links but those I wished to remove. The Tinderbox *Roadmap* view, which shows all links to/from a node and all its aliases, proved useful in this task. Since I started with more than 60 nodes, some node aliases were also removed. Had I culled the original spatial map of linked nodes, the removal of a node or link would have been final, since Tinderbox does not store versioning information. Instead, aliases enabled me to try multiple arrangements of the same data.

Two very different link-viewing behaviours proved equally useful in the design of the geodesic hypertext. In some cases, aliases needed to retain their unique properties. The spatial map needed to display only unique alias links. However, the *Roadmap* view, also an essential part of the planning process, does show all links to/from both the original and its aliases.

A more formal consideration of issues relating to transclusion and aliases in spatial hypertext should be made. However, the designers of hypertext tools are encouraged to give users flexibility in displaying and sorting lists of the properties and links of original nodes and their transclusions.

4.2 Link Culling in Gzz

For the purpose of this case study, both the design methods and final geodesic hyperstructure were tested on Gzz 0.6.3. This software is an implementation of Nelson's zzstructure convention "for data and computing." Zzstructure and its related tools intend to be "a generalized representation for all data and a new set of mechanisms for all computing" [10]. Gzz presents all data as cells, which may be linked orthogonally along a potentially infinite number of user-definable dimensions. Links are directional, and the users may select one of several views to navigate among cells by following the links. Zzstructure-based software is *not* hypertext software, but rather software that visualizes a data structure that is recommended as a base for future hypertext applitudes, the equivalent of applications in the ZigZag* system[10].

Gzz would not have easily aided the task of link culling. Cells are limited to two links per dimension, and most of the views only display a maximum of three dimensions[10]. Thus, an initial state of many interlinked cells, a necessary starting point for link culling, would be difficult to create. Although zzstructure clones exhibit similar tendencies as Tinderbox aliases, other Gzz behaviours make link culling difficult. Cursors only navigate along links in the currently-viewed dimensions. Collections of cells can disappear from view when nodes are unlinked. Worse, navigation within the hypertext can become more difficult every time a link or cell is removed.

Difficulties with one interface to ZigZag do not indicate problems with the underlying storage mechanism. One can imagine an applitude designed to aid link culling techniques. Even in Gzz, navigation problems could be partially alleviated through the use of multiple view windows and a clone-based authoring strategy modeled after my Tinderbox alias methods

4.3 Case Study in Geometric Hyperstructure

Philadelphia Fullerine provides a useful case study in issues related to geometric hyperstructure. In an overview of hypertext narrative patterns, Bernstein discusses early speculation on carefully-structured hypertexts. This research suggested that symmetrical and three-dimensional hypertexts would reduce user disorientation. Bernstein also points out that other research suggests that "clearly marked paths [rather than uniform structures] help keep readers oriented" [5]. Rather than discuss the best way to promote ease of reader navigation, Bernstein brings together later research to show that the issues of disorientation may not be a source of concern. Disorientation, particularly in literary works, is much less detrimental than initially predicted. According to Bernstein, such disorientation may be positive [5].

Personal experience with *Philadelphia Fullerine* and initial reader response have indicated that a symmetric, three-dimensional hyperstructure does increase user disorientation. Rather than understand the location of individual items in relation to the whole, readers use a contextual navigation strategy similar to the strategies used in the cursor-centric view of Gzz[10]. However, navigation of the geodesic structure is difficult to analyze electronically, since attempts to compare software navigation strategies to reader experience were hindered by the difficulty of properly entering a geodesic hyperstructure into Tinderbox and Gzz.

4.3.1 Storing Geometric Hyperstructure

Neither Tinderbox nor Gzz are able to properly represent the geodesic structure of Philadelphia Fullerine. The reason is simple. Both software tools rely solely on directional links while the sculpture's links have no direction. This disparity may not be a major problem, but it caused some mild inconveniences in Tinderbox. Storing and viewing the structure in Gzz was more difficult; the process required modification both to the geodesic structure and eventually to the basic operation of Gzz.

4.3.1.1 Link Directionality in Tinderbox

In Tinderbox, I made directional links for the hypertext structure as was convenient. Implicit linking, which is based on adjacency, more closely resembles the actual link strategy of the sculpture, but link culling cannot be easily done for implicit structure.

* ZigZag is a registered trademark in the USA for the zzstructure-based software of Project Xanadu.

Furthermore two-dimensional Tinderbox maps cannot show implicit three-dimensional structure[7].

The unpredictable directionality of explicit links in Tinderbox caused some essential links to be absent in views that only showed outward links. In other views, the three links would sometimes be displayed separate columns. However, Tinderbox, with its multiple ways of looking at data, does not lose much by solely allowing directional links. Indeed, while Bernstein notes the theoretical usefulness of bidirectional links, he suggests that their absence in systems like Storyspace and Tinderbox is not a great loss[4].

4.3.1.2 Storing Geodesic Hyperstructure in Gzz

Entering a geodesic hyperstructure into Gzz proved a much greater challenge, since zzstructures are by nature orthogonal, and geodesic geometry is not. Furthermore, the posward or negward direction of zzstructure links often determines the arrangement of cells on the screen and the commands used to traverse them[10]. Gzz views of the hyperstructure provided some surprising results, in some cases introducing an unexpected hierarchy to the structure. However, after some modifications to the Gzz dimensions, the Gzz compass view closely modeled navigation of the physical sphere.

Many different zzstructures could have been used to store the geodesic hypertext -- as many as the possible combinations of 140 directional links in a geodesic pattern along a potential of up to 140 different dimensions. However, a precise representation of the spatial or conceptual structure is not possible. Other hyperstructures, like a hypertorus or a hypercube, can easily be stored in a zzstructure, since both the structures and the software are orthogonally-based. On a sphere of triangles with full rotation, links occur in many spatial directions.

One might next seek a conceptual equivalent of the structure. This is also difficult to model. On the geodesic sphere, no two links occur in the same direction. Furthermore, each triangle is a potential center, origin, and end. All links in the sphere are outward. Since zzstructure links must have both a posward and negward side, it is impossible to encode the geodesic structure. A zzstructure only allows two links in a single dimension, but a conceptual geodesic structure requires three.

I could have used a unique dimension for each unique link, but this introduces further problems. Traversal becomes difficult, since the direction of links determines both the display and traversal of links. Visualization and navigation become unpredictable.

As a compromise, I used two dimensions. This biases link traversal to the dimension which contains more links. The first dimension, called "clockwise," connects loops (ringranks) of 5 cells. Following these ringranks in a posward or negward direction would be like rotating the sphere around the vertices that form the center of pentagonal triangle groupings. Twelve of these ringranks were used. I created a second dimension to join the ringranks. These links would simulate the act of following a link across the imaginary pentagonal borders.

Directional links worked well along the ringranks, but traversing the other dimension was difficult. Directional links caused uncertainty about which direction (and thus which key command) would be needed to traverse along this dimension. To ease the difficulties traversing the structure, the trans-rinrank links were made in both posward and negward directions. Thus, each triangle used four links to link to three other cells. This strategy resolved the interface navigation difficulties for many Gzz views.

Gzz limits visible nodes to those close to the current cursor-selected cell. The three-dimensional interface of Gzz centers the view on the currently-selected cell, much like viewers of the physical sculpture orient individual triangles to face themselves . Thus, I expected the Gzz views to duplicate elements of the three-dimensional physical experience more closely than other software. The results were unexpected.

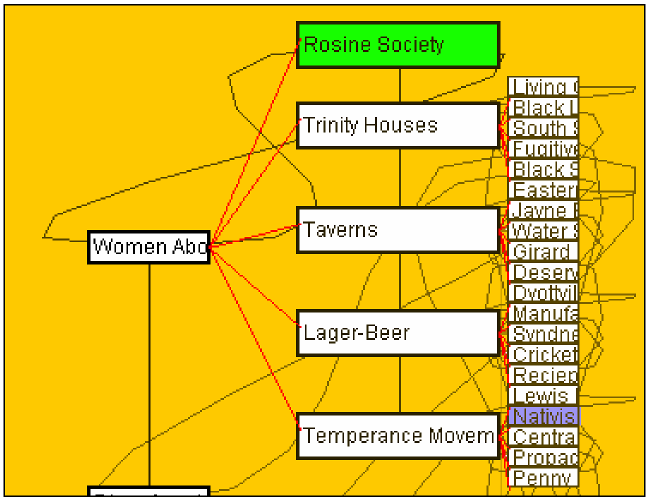

Figure 4. HTree View of Geodesic Structure

I expected hierarchical views to display the hypertext as an infinitely-symmetrical, infinitely-balanced, recursive tree, with each node pointing to three children. Instead, Gzz showed a balanced tree with five children per parent (see figure 4). The *MindSundew* view also replicates a five child hierarchy (figure 5).

This unexpected behaviour demonstrates ZigZag's definition of hierarchy. In the *HTreeView* shown in figure 4, parent/child relationships are shown on the x axis. However, since ZigZag only allows one link in any direction, only one of the children is linked directly to the parent. Children are defined as those cells which are in the same rank as a cell which is linked directly to the parent along another dimension. Hierarchy indicator lines are shown by the *HTreeView* but not traversable or encoded in the zzstructure. A potentially infinite number of cells and links may lie between a cell and its parent in a hierarchy. These cells must be traversed to reach the parent. Gzz interprets the interlinked , five-cell ringranks as a recursive hierarchy with each cell as parent to five other cells, even when each cell only links directly to three other user-defined cells.

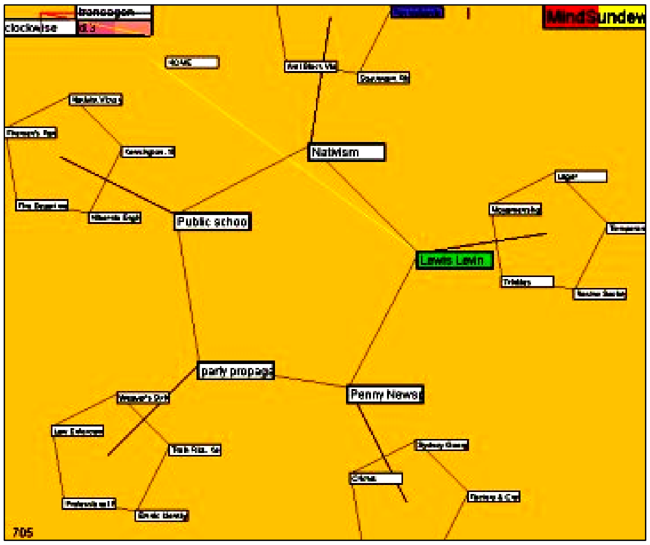

Figure 5. MindSundew View of Geodesic Structure

Although the *MindSundew* view (see figure 5) seems to most closely model the geodesic shape of the sculpture, the difficulties with link direction do blemish the graphical elegance of this view. Unless the trans-ringrank links are duplicated in both directions, the view shows a structure which seems to be collapsing upon itself, obscuring some of the options. Furthermore, the *MindSundew* view organizes individual cells into implicitly-linked groups of 5 cells.

To properly display the geodesic hypertext, a view should not introduce or suggest more links than are explicitly made. It also needs to show links in two or three dimensions without emphasizing links made in any single dimension or rank. The view which most closely matches this is the *Compass* view.

The *Compass* view centers the currently-selected cell, surrounded by all connected cells in all dimensions. As seen in figure 6, the *Compass* view for the geodesic structure displays links to two extra cells. This is because Gzz automatically records a *cellcreation* dimension, along which all nodes are linked in the order of their creation. The *Compass* view naturally displays these links. By disabling this dimension by unlinking it from the dimension list, the desired effect was achieved (see Figure 7).

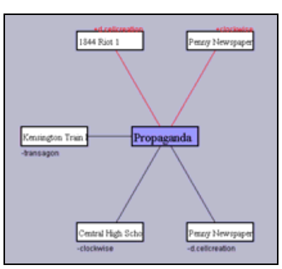

Figure 6. Compass view

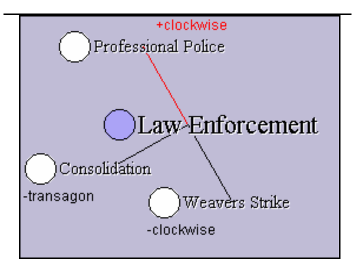

Figure 7. Compass view after unlinking a dimension

In the *Compass* view of the geodesic structure, only four cells are visible at any one time. To navigate the structure, users select one of the adjacent cells. The view will then re-center on the newly-selected cell and display the three topics adjacent to this new cell. This is the closest electronic representation of navigation of the geodesic structure I have so far achieved.

Although some enthusiastic supporters of zzstructure suggest that it may be easily capable of storing any type of data structure, this geodesic hyperstructure demonstrates some limitations of the model, limitations which have useful effects and are part of ZigZag's conscious, intended design**Error! Reference source not found.**[9][10]. However, most difficulties with current ZigZag implementations can be overcome through the design of custom cell conglomerates and applitudes.

4.3.2 Navigating Goemetric Hyperstructure

Visually finding specific nodes is difficult in all representations of this hyperstructure: the sculpture, the Tinderbox spatial map, and the Gzz implementation. Non-orthogonal structures allow fewer navigation strategies than structures which link along the axes. This compounds the navigation difficulties common to many geometric hyperstructures.

Parunak's overview of hypermedia navigation identifies a number of navigation strategies for geometric hyperstructures. Specific nodes are often found by remembering a list of turns and distances. Parunak calls these strategies the "path and distance" navigation strategies. Parunak also notes that other key features aid navigation in symmetrical structures: "without a distinguished point, a hypertorus (like a sphere) does not support either[texture or comparability], even if it is ordered"[13].

A geodesic hyperstructure presents fewer navigational aids than an orthogonal structure. Distance path methods require at least two links in the same direction. Even in a hypertorus, many links can occur in sequence in the same direction. In the geodesic sphere, each link requires a turn. Only texture-based strategies -- which are weak in this hypermedia-- and path-based strategies may be used in a geodesic structure. To find specific items on Philadelphia Fullerine, users are more likely to wander aimlessly until they spot the item they desire than they are to plot the path from disparate nodes. The same is true for even the author.

A hypertorus may suffer from many of the same navigational problems as a sphere. Without a clear starting point, distance and path strategies lessen in usefulness. A mental understanding of the entire structure becomes important. To plot a path, one would need to imagine the relative locations of the current node, the target node, and the nodes in-between.

This disorientation is expected and allowed in link-traversing, current-node-centric views for systems like the Gzz *Compass* view. Such systems constantly change visual representation and geometry; they concern themselves with clearly showing the immediate context more than the entire hyperstructure. However, they suffer the same navigational difficulties as the sculpture, since path and distance navigation requires either a fixed starting point or a thorough mental understanding of the entire structure.

A hypercube simplifies navigation in several ways. First, implicit groupings of links on the sides of the cube exemplify what Parunak calls "texture." Furthermore, the addressing schemes suggested by Parunak are also useful in a hypercube. The six sides of a cube provide six areas on which the location of a single node is fixed. Whenever a node is centered in a sphere or hypertorus, the locations of all other nodes shift. Viewed in the subdivision of a single side of a cube, the node locations remain constant.

Even in the spatial map view of Tinderbox, the uniform geodesic structure can be difficult to navigate. The task of representing a three-dimensional structure in two-dimensional space is the least difficulty. In fact, a two-dimensional map aids navigation by fixing the position of individual nodes. The primary difficulty in my arrangement of Tinderbox data is the lack of texture; large number of uniform nodes are only distinguishable by their text. For new viewers, this is very disorienting. However, one is able to find specific nodes much more quickly within Tinderbox than within Gzz or on the actual sphere. With both Gzz and the sphere, only nodes in the immediate context are displayed. In Tinderbox, they are all displayed. One could also improve navigation in any of these systems by adding texture. In a move similar to placing nodes on the side of a cube, regions could be color-coded.

Practical uses exist for the disorientation caused by a spherical geodesic structure. Indeed, I do not intend users to seek specific information. This hypermedia does not intend to present the same information at each viewing, but rather to provide a unique experience for every viewer each time it is used. Had I wanted viewers to find information more easily, I could have followed my original intent to make a cube-shaped sculpture.

In fact, a significant number of *Philadelphia Fullerine's* viewers choose to abandon link-based navigation. Although the links offer them a linear way to progress through the content, viewers enjoy spinning the ball. They will often seem to pick nodes at random, based on the images that appeal to them. Since nodes are interlinked but not dependent on their adjacent nodes, this strategy provides a satisfactory alternative user experience. For a broader range of purposes, both contextual and complete viewing of hyperstructure can be useful. In some cases, users wish to find specific information. In other cases, they may wish to explore. Contextual views can helpfully limit and guide users, but they can also complicate navigation. Views which display many nodes can free users, but they can also reduce a reliance on links to establish narrative and conceptual flow.

5. CONCLUSION

Philadelphia Fullerine is a unique physical manifestation of hypermedia. This creative nonfiction hypermedia presents a useful case study for managing a nonfiction hypertext project within hypertext software. It also presents a case study in the creation of a literary hypermedia through link culling with spatial hypertext.

The issues raised by attempts to store and interact with this hypermedia within hypertext software raise questions which I hope will influence more thoughtful discussion about such systems. I also hope to see more consideration of the relationship of spatial hypertext to transclusion and further discussion on the proper storage and visualization of complex hyperstructures.

The choices of user interface and hyperstructure made by software designers and authors define users' experience with a hypertext tool or work of art. It is my hope that creators of future hypertext software will be further encouraged to provide a wide variety of user interface options to enable the kind of flexible use I needed to design, create, and present the *Philadelphia Fullerine*.

6. ACKNOWLEDGMENTS

Elizabethtown College partially funded *Philadelphia Fullerine*. Thanks to Mark Bernstein and Jim Rosenberg for clarifying the term "sculptural hypertext".

Thanks to Tuuka Hastrp of Jyväskylä University, for suggestions for storing the geodesic structure in Gzz.

REFERENCES

[1] Mark Bernstein, *Card shark and thespis: exotic tools for hypertext narrative*, In *Proceedings of the twelfth ACM conference on Hypertext and Hypermedia*, August 14-18, 2001, Århus, Denmark

[2] Mark Bernstein, *Collage, Composites, Construction*, In *Proceedings of the fourtheenth ACM conference on Hypertext and Hypermedia*, 2003. Nottingham, U.K.

[3] Mark Bernstein. *Hypertext and the linearity of history*. HypertextNow. Watertown, MA: Eastgate Systems. <http://www.eastgate. com/HypertextNow /archives/History.html>

[4] Mark Bernstein, *Storyspace 1*, In *Proceedings of the thirteenth conference on Hypertext and hypermedia*, June 11-15, 2002, College Park, Maryland, USA

[5] Mark Bernstein, *Structural patterns and hypertext rhetoric*, ACM Computing Surveys (CSUR), v.31 n.4es, Dec. 1999

[6] William Gillespie, *Letter to Linus: A Hypercube*. Spineless books, 2002. <http://www.spinelessbooks.com/lettertolinus/index.html >

[7] Catherine C. Marshall , Frank M. Shipman, III, *Searching for the missing link: discovering implicit structure in spatial hypertext*, In *Proceedings of the fifth ACM conference on Hypertext*, p.217-230, November 14-18, 1993, Seattle, Washington, United States

[8] Catherine C. Marshall , Frank M. Shipman, III, *Spatial hypertext and the practice of information triage*, In *Proceedings of the eighth ACM conference on Hypertext,* p.124-133, April 06-11, 1997, Southampton, United Kingdom

[9] Adam Moore , James Goulding , Tim Brailsford , Helen Ashman, *Practical applitudes: case studies of applications of the ZigZag hypermedia system*, In *Proceedings of the fifteenth ACM conference on Hypertext & Hypermedia*, August 09-13, 2004, Santa Cruz, CA, USA

[10] Theodor Holm Nelson, *A Cosmology for a Different Computer Universe: Data Model, Mechanisms, Virtual Machine and Visualization Infrastructure,* In *Journal of Digital Information*, v.5.1 n.298, June 16, 2004

[11] Theodor Holm Nelson. *Literary Machines*. Mindful Press: Sausalito, 1981. Reprinted by Eastgate Systems.

[12] J. Nathan Matias. *An Accordion for the World*. In Tekka, vol. 7. Dec., 2004 <http://www.tekka.net/07/?Accordion>

[13] H. Van Dyke Parunak, *Hypermedia topologies and user navigation*, In *Proceedings of the second annual ACM conference on Hypertext*, p.43-50, November 1989, Pittsburgh, Pennsylvania, United States