

Browsers with Changing Parts: a Catalog Explorer for Philip Glass' Website

Harry Hochheiser¹

Human-Computer Interaction Lab &
Department of Computer Science
University of Maryland
College Park, MD 20742
+1 301 405 2725
hsh@cs.umd.edu

¹ Work done at IBM T.J. Watson Research Labs, Hawthorne, NY

ABSTRACT

The development of navigational tools for a web site devoted to a catalog of musical compositions offers a variety of design challenges. A combination of techniques developed from information visualization research - including starfield displays, dynamic queries, and zoomable user interfaces (ZUIs) - was used to construct a prototype browser for possible use in a web site dedicated to the work of Philip Glass. After a discussion of the goals and potential users of the site, this paper describes several prototypes that were developed and how they informed the design of a zoomable starfield browser. Unresolved design challenges and possibilities for future work are also discussed.

Keywords

Interactive Visualization, Music, Biography, Catalog Browsing, Zoomable User Interfaces (ZUIs)

1. INTRODUCTION

The design of web-based tools for browsing catalogs of musical works involves numerous problems. These tools require the integration of music into the primarily visual environment of the Web, in a manner that supports users of varying musical backgrounds in pursuit of a wide range of goals. Information visualization techniques such as dynamic queries, semantic zooming, and starfield displays may be useful, but application of these techniques to musical data has been limited, and their utility for tasks involving significant non-visual components is unknown.

These concerns and others arose during investigations of possible designs for browsing and navigation tools to be used

his operas (*Einstein on the Beach*), film scores (*Kundun*), symphonies and other works, Glass has a catalog of several hundred compositions.

Recently, Glass and his production staff have been exploring possible designs and functionality for a Philip Glass web site. In addition to serving as a web presence for Philip Glass, this site would provide educational materials and promote Glass recordings, performances and other products. Another, more innovative proposed application of the site involves licensing. Glass currently licenses his music to a wide range of performers, filmmakers, and other artists. The web site is seen as a potential tool for supporting the identification of works to be licensed, and for completion of the eventual transactions.

In collaboration with Glass' staff and the ECommerce Intelligence Group of IBM's T.J. Watson Research Labs, we defined potential users and user tasks and developed prototype components for this site. Although primarily focused on the design and development of components that could be used in a production web site, these efforts also provided an opportunity to address some of the challenges involved in a novel application of information visualization techniques. This paper describes several prototypes that were developed, and their influence on the design of GlassEye: an interactive visualization tool that combines musical playback with a starfield display [1] in a Zoomable User Interface [4].

2. BIOGRAPHICAL WEB SITES

Our exploration of possible designs for the Glass site began with an informal survey of web sites devoted to musicians and composers. As there are hundreds (if not thousands) of such sites, this effort was not intended to be comprehensive. Rather, we hoped to understand the range of currently existing sites in a manner that might help frame our discussion of the goals of the Glass site. To that end, we examined several existing sites for both classical and pop musicians, including both officially-sanctioned and "fan" sites. This survey also included several "unofficial" Philip Glass sites. This exploration led to a set of categories that describe the existing sites:

within a web site devoted to the music of Philip Glass. Known for

Encyclopedic sites present information in a reference format designed (hopefully) for ease of search and retrieval, usually for use in answering specific questions. Examples include lists of an artist's work, timelines, lists of songs by album/year, list of performances, etc. Appropriate navigation tools for encyclopedic sites may include searches, indices, and hierarchical menu traversal.

Commercial sites are strongly oriented towards selling merchandise. Navigation through a catalog of items available for purchase might involve browsing through a store-like setup or possibly browsing content pages with interspersed commercial/product links.

Educational sites contain traditional educational material, designed to teach users about the musician. Usually, they include images and linear explanatory text with minimal links. Timelines might be particularly useful for these functions.

Fan-oriented sites often contain pictures, links to other sites, personal comments, and similar "tribute"-like information. Random browsing and exploration are usually appropriate navigation techniques for these sites, particularly if these pages are used as "hooks" aimed at enticing users into further, deeper exploration of a site.

Contextual sites provide biographical information that places an individual in a historical, social, and artistic context. For example, a site dedicated to contemporary music might have a single page on Glass, along with pages describing other modern composers. This contextual information could be provided through judicious use of side links and sidebars.

These explorations did not narrow the range of goals for the Glass site: instead, it became clear that the ideal site would fill *all* of these roles. However, this characterization did help focus discussions of the impact of various goals upon the eventual site design. Specifically, we realized that the open-ended browsing characteristic of fan-based sites might not be compatible with the presentation of content needed for a commerce site. On the other hand, the intrusion of commercial content might detract from the experience of exploring educational material.

3. REQUIREMENTS, GOALS, AND TASKS

Based on our discussion with Glass' staff and our examinations of related web sites, we developed an informal set of requirements for the site, along with a framework for describing users tasks. These requirements provided a challenging and often inconsistent set of goals for the site.

3.1 Functional Requirements

E-Commerce: Although CD, book, and video sales would be supported, the main E-commerce function of the Glass site would be to arrange for licensing of recordings and rights. Support for these goals would be provided by an electronic version of Glass' catalog, along with tools that potential licensees could use to search the catalog for appropriate pieces. Relevant information

would include instrumentation; length of pieces; background details about the origin, themes, and context of the works; and (if available) descriptions of tempo, mood, and other criteria that may be provided by Glass's production team.

Education: The Glass site might provide a wide range of possible educational functions. Facilities might provide support for users unfamiliar with Glass's work and interested in learning more, knowledgeable fans looking for specific details, students working on academic studies related to Glass and his music, and others. Guided tours, frequently asked questions (FAQs), and suggestions ("if you want to find...") might be used to help orient new users. This introductory material might also create of an inviting and accessible site and provide additional navigation support for more advanced users.

Design of educational materials for a wide range of users presents significant navigational challenges: a properly designed system should provide experts with the facilities needed to identify necessary details, while providing novices with an approachable and accessible introduction to the material. One way to achieve that goal might be a "layered" interface, which would enable users to progressively reveal greater levels of detail.

Exploration: Some users of a Glass web site are likely to come to the site with goals that are incomplete or otherwise not well specified. For such users, developing a better understanding of their goals may be a desirable end unto itself. Browsing facilities might encourage users to explore while providing cues that help users avoid disorientation. In addition to helping users clarify goals, these tools might add to user satisfaction by providing the enjoyable experience of "discovering" data on the site. Exploratory components might combine e-commerce and educational functionality, but with an added emphasis on interactive facilities. Possibilities include musical scrollbars, musical construction kits ("build your own Philip Glass work"), and facilities for identification of appropriate pieces for licensing. Implementation of such facilities in a web context may be limited by browser capabilities.

Searching: Users interested in licensing Glass music for use in a particular context may be interested in identifying music that meets appropriate subjective criteria. In addition to searching for a piano piece or a chamber work, these users may be interested in music that is particularly intense, fast, or triumphant. Support for searching along these lines may be appealing, but generating the necessary descriptions of the pieces in the Glass catalog may be labor-intensive. However, when combined with tools that support exploration and readily available audio samples, subjective searches could provide a powerful tool for generating licensing arrangements.

3.2 Aesthetic and Philosophical Requirements

Generality: While e-commerce is a primary function of the site, users should be able to enjoy and learn from the site without necessarily making a purchase.

Inclusion of Visual Arts: Visually, the site should incorporate aspects that reflect Glass' collaboration with visual artists in writing music for dance, film, and stage.

Novelty: Ordinarily, use of new technologies strictly for their novelty value is not encouraged. However, the Glass site should be in harmony with Glass' reputation as an innovator. The site should be "cool", and should strive to go past standard linear forms often found in web designs. A site that makes constructive use of new technologies to support users in their exploration of the Glass catalog would ideally generate an enthusiastic user community while extending the range of Glass' innovations into a new realm.

Distinctive Identity: Glass' music has a distinctive, recognizable style. The Glass site should be similarly distinctive: Ideally, the site will create a "brand" which will be easily recognized as belonging to Philip Glass. This might be achieved through appropriately designed icons and backgrounds, images from stage productions of Glass works or movies, pictures of the artist, or other visual design techniques.

Focus on the music - Access and Understanding: The web is still a primarily visual medium. Streaming audio and MP3 provide audio content that is often completely disconnected from the visual content of web sites. Despite these trends, the site should emphasize Glass' music as primary. Site layout, navigation tools, Java applets, and other features should serve the goal of connecting users to the music, in order to generate (or extend) their understanding for and appreciation of Glass' works.

3.3 User Tasks

As expected, navigation of the Glass catalog and selection of items of interest were identified as the main user tasks that must be supported. This simple classification was augmented by the definition of a series of dimensions that could be used to describe user tasks:

- User *expertise*: how much does the user know about Glass and his music?
- User *perspective*: is the user a fan? A student? A musician? All of the above?
- Level of *focus*: is the user engaging in free-form exploration, or searching for answers to highly specific questions?
- What is the user's *goal*? Finding information? Buying a CD?

These categories were used to develop more detailed lists of potential user perspectives and tasks, which then influenced our discussion of design alternatives. In particular, this analysis led to the explicit decision that a single interface supporting all of these tasks would be ideal.

4. PROTOTYPES

Development of a series of prototype components proceeded in parallel with the site exploration and system requirements work described above. These prototypes provided perspective on ideas

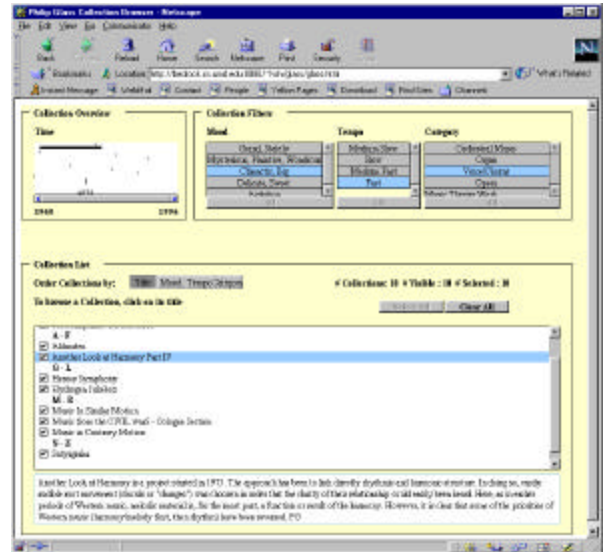


Figure 1: A collection browser for the Glass Catalog



suggested by Glass' staff and informed the development of the GlassEye browser.

4.1 Timeline Browsing

A chronological view of the catalog as a history of Glass' career led us to consider the application of existing tools designed for browsing of temporal data. LifeLines [19] and the National Digital Library collection browser [21], both developed at the University of Maryland's Human-Computer Interaction Lab (HCIL), provide timeline displays to facilitate browsing data sets.

Figure 2: A prototype home page for the opera Akhnaten.

Both Lifelines and the collection browser support dynamic queries – “interactive user control of visual query parameters that generate a rapid (100 msec update) animated visual display of database search results”[23]. The collection browser, originally developed for the Library of Congress’ American Memory project, includes simultaneous menus [12] for selection of categorical information, along with tightly coupled displays of title and summary information.

The collection browser was adapted to display a subset of the detail from the Glass catalog, using subjective measures of mood and tempo, along with the objective categorization of the type of work (opera, orchestral music, etc.). These attributes were used as input to a series of simultaneous, dynamic query [23] menus that can be used to quickly add and remove works from the list being viewed.

The combination of a linear, categorized list of compositions with a temporal view of the catalog facilitates browsing through the data. Tightly coupled highlighting of selected pieces helps users identify correlations between the timelines, categories, and the linear list.

4.2 Sample Web Pages

A fully developed Glass site might contain individual HTML pages describing each of the works in the catalog. A prototype home page for the opera *Akhnaten* was developed to illustrate potential content and layout possibilities (Figure 2). At the top of the page, navigational links providing access to information about related works, products, commentary, performances, and audio clips are centered on an image of a Tibetan mandala. Further down the page, summary texts, excerpts from liner notes, lists of collaborators, and other details are given, along with links that can be used to drill down into the site for greater detail. The page ends with global navigation links aimed at helping users find their way to other appropriate areas on the site. The mandala at the top of the page was chosen based on Glass’ interest in Buddhism (and his work on the soundtrack for *Kundun*), and might serve as a navigational theme that could be used throughout the Glass site.

4.3 Musical Scrollbar

Inspired by a suggestion from Glass’ staff, the Musical Scrollbar was designed as an interactive browser that combines display of detailed information about various Glass pieces with audio playback of selected segments. Implemented as a Java applet, the musical scrollbar contains detail fields describing the name of the piece, the work from which it was taken, the year, and the premier location and date. Additional text areas provide background information, instrumentation details, and comments.

The applet is controlled through a control at the top of the window. While the name “scrollbar” has stuck, the control is actually implemented as a slider: users can drag the thumb across the horizontal range, or they can click on the boxes at either end to move left or right by one selection. As the user drags, the detail

display areas are updated dynamically, based on the position of the slider.

There are two possible modes for the musical playback. If the user leaves the thumb on a given piece, or if the increment/decrement boxes at the end of the slider are chosen, a short audio clip from the selected piece is played. If the user drags the thumb, a special “dragging music” selection will be played until the user releases the thumb, at which point the selected piece will begin playing. Audio playback can be muted through selection of a checkbox provided on the applet.

This demonstration presents a basic model of a system that can be deployed using current browser technology. Extensions such as multiple scrollbars for scrolling in other dimensions might be used to provide additional navigational power.

4.4 Zoomable Space

Zoomable User Interfaces (ZUIs) are 2.5 D spaces that use magnification to manage object representations that differ with scale. ZUIs support the use of “semantic zooming” – the revelation of progressively greater levels of detail with increased magnification [4] – for the construction of overview and detail displays, in which users can access information on topics of interest by magnifying appropriate areas of the zoomable space.

As a means for exploring the application of zooming to the Glass catalog, we implemented a zoomable visualization of selected elements from the Glass catalog. Implemented with Pad++[4], this visualization is based on a layered timeline: compositions are displayed in chronological order in the horizontal direction. Vertical stratification is used to separate the pieces by type: operas on the bottom, symphonies in the middle, chamber music above, etc. Graphical representations of the compositions are scattered throughout the 2-dimensional display space, in effect creating a starfield display [1].

Initially, a 30-year timeline displays an overview of Glass’ career. Users can manually zoom or click on a hyperlink for any of the pieces to change the scale to present details about a given piece, including information regarding recordings, premieres, and liner notes. Zooming further, users might be presented with audio samples or more extended textual descriptions. At any stage, hyperlinks are available to return users to previous levels of detail.

The example Pad++ visualization developed for the Glass catalog also includes menus arranged around the perimeter of the zoomable space. These menus provide additional dynamic query filtering mechanisms that might be used to identify music by type, popularity, instrumentation, or other criteria. While not functional in the Pad++ prototype, these menus provide a hint of how dynamic queries might be combined with a zoomable visualization.

4.5 Lessons Learned

The prototype web pages provided a vivid demonstration of a difficulty that is likely to be faced by any system for browsing a

complex data set such as that of the Glass compositions. Each of the works in the catalog has a wide range of auxiliary data, including performance details, product details, explanatory texts, instrumentation information, and other data that must be navigated.

Furthermore, larger compositions such as operas might be considered as collections of individual pieces, each of which might have a related data set. This clearly presents a challenge for interaction tools: in addition to supporting navigation *to* each of the components in the catalog, these tools must also provide assistance for navigating *within* the individual pieces.

Several of the prototypes described above were shown to members of Glass' production staff, in order to generate discussion and comments on the various designs. Since the site is likely to present a very public image of Glass and his work, we felt that the success of any system would be contingent upon the approval of Glass and his staff.

Two of the prototypes – the collection browser and the Pad++ browser - generated particularly enthusiastic responses. The view of the catalog provided by the collection browsers' tight-coupling between the catalog data and the timeline was particularly attractive to Glass' staff, as it supported highly interactive browsing of the catalog. They also found the zooming and the large browsable space of the Pad++ prototype to be both intriguing and fun. Based on these reactions, and on our interest in exploring the possibilities of building a zoomable starfield display, we decided to build a new prototype based on a combination of the temporal display of catalog data in the collection browser and semantic zooming in the Pad++ prototype.

5. GLASSEYE PROTOTYPE

Building on the Pad++ prototype, the GlassEye browser combines a starfield display for visualization of the catalog

browser with a zoomable interface that provides greater detail at higher levels of magnification. The browser is implemented using Java 1.2 and the Swing interface toolkits. Zooming support is provided by Jazz [5], a descendant of Pad++ written in Java.

In the spirit of continuous prototyping, GlassEye was designed to provide a vision of what a fully operational browser might look like. As such, it demonstrates possibilities, and does not address all of the issues above. The use of Jazz also raises significant issues for eventual deployment: the networking and processing overhead associated with the GlassEye code and the Jazz libraries would likely render the applet unusable for users with slow connections or slower processors. As a result, this prototype should be viewed as a demonstration of concepts, and not as a production system.

In the initial GlassEye display, the horizontal axis is used as a timeline displaying the year of composition. The vertical axis is used to display a rough count of the number of instruments required to perform a piece, thus providing users with one approximate perspective on the complexity of the piece – in Glass' case, ranging from solo piano to orchestral pieces. This initial display was chosen as a representative guess as to which views might be most useful for users: issues regarding semantics of starfield axes in this context will be discussed below.

Within the starfield display, each composition has a single representation. For titled pieces, the name of the work is used. For untitled pieces, a circle is used. In either case, color-coding is used to display the type of the piece (orchestral, stage, or chamber music). Size coding based on an arbitrary "popularity" measure is used to emphasize works that users are most likely to find interesting. The resulting display provides a single-screen overview of the entire catalog.

Dynamic query filters [23] are provided for easy filtering based on additional attributes: checkboxes allow for selection of pieces by their type, while sliders for popularity, length, and other attributes line the bottom of the display.

The browser also provides a smaller, secondary window containing an overview of the starfield space (Figure 3). For assistance in navigation, this window provides a view that remains constant even as the magnification in the main window changes.

5.1 Zooming for Detail

GlassEye uses zooming to make detailed information about each piece available at higher levels of magnification. The initial implementation provides two levels of detail about each piece. First, background information describing the work, collaborators,

premiere date, related works, are provide, possibly along with one or more images. At the next level, extended excerpts from liner notes or other further details are given (Figure 4).

Magnification hyperlinks provide the simplest approach to zooming in the GlassEye browser. Each composition is a hyperlink: clicking on the textual or iconic representation of a work will cause an animated magnification, resulting in the appropriately magnified display of detail about the selected piece in the center of the screen. All displays that are targets of hyperlinks include links that can be used to return to the previous magnification. Currently, these links are depicted as up-arrows, indicating the upward motion of the viewpoint that would lead to a decreased level of magnification.

Hyperlinks in the GlassEye browser are not limited to change


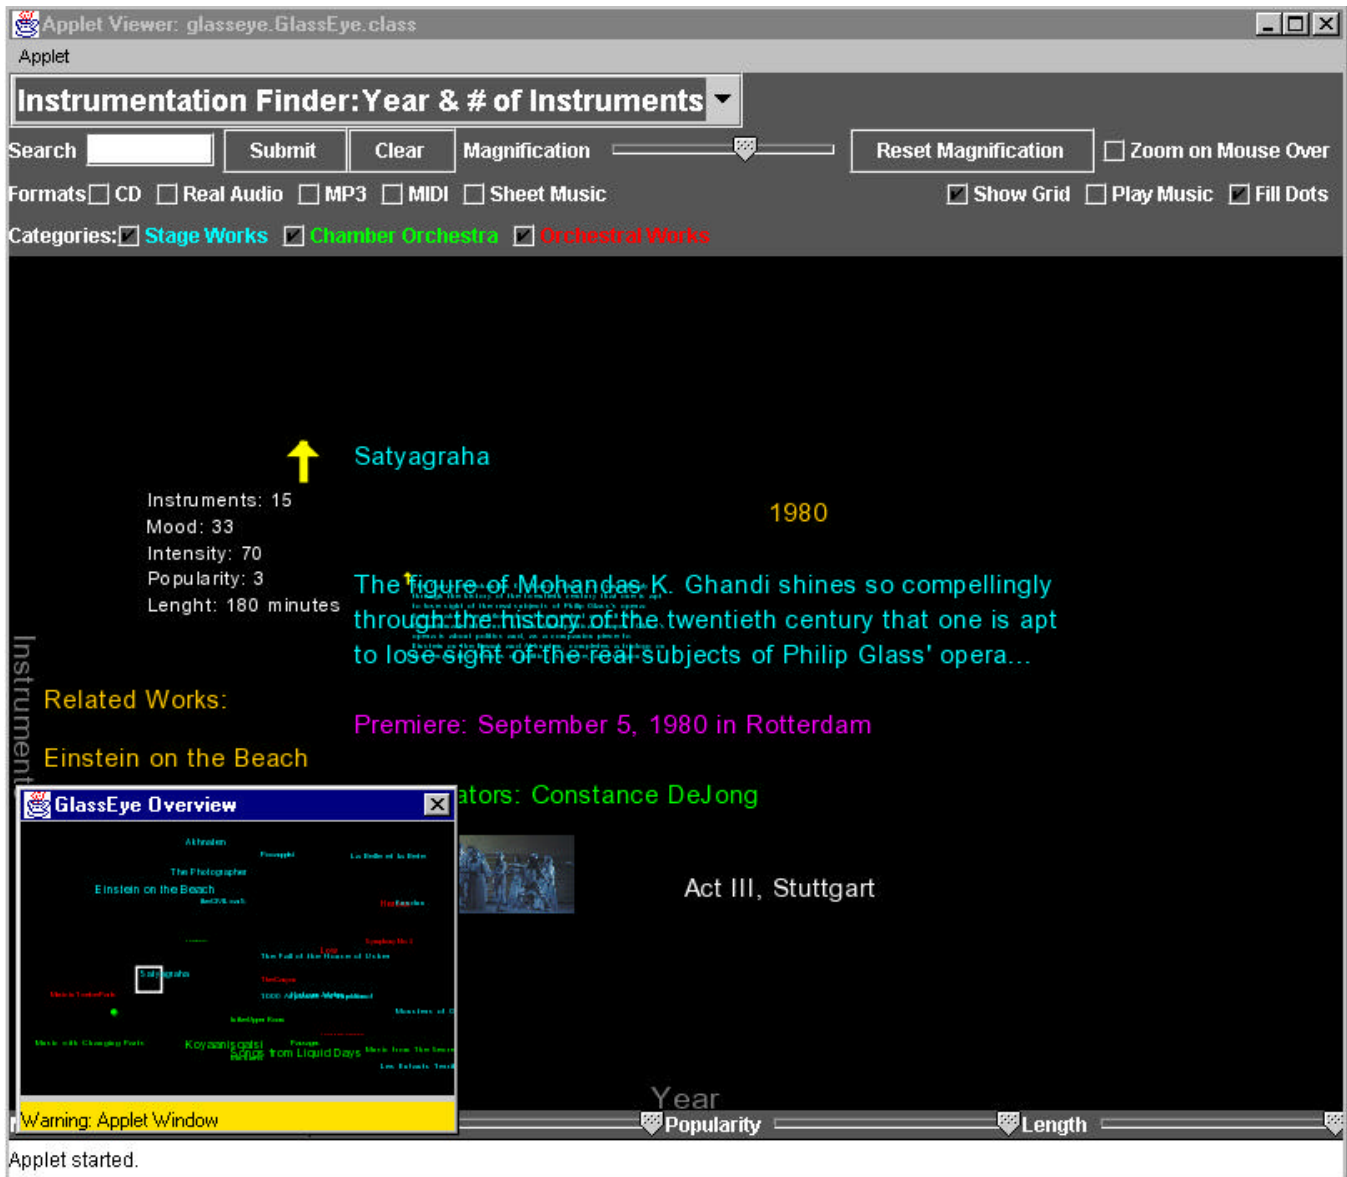


Figure 3: GlassEye browser with overview window



of scale: translation links are used to connect related compositions. Links are presented to users as outline borders that appear around linked objects on mouse-over. At the main level, further emphasis on links is provided by a changed color for the object that has the current focus.

Zooming is also used as a navigational tool. As GlassEye's starfield display can become somewhat crowded, magnification of areas of interest might help users distinguish between labels that are located near each other.

Furthermore, semantic zooming is used to bring individual detail items into view as magnification levels are increased, providing some residue or "scent" that might guide users in their continued navigation [11].

Users can choose between two possible approaches for control of general magnification. Clicking on the background of the display adjusts the magnification, with the left mouse button increasing the magnification, while the right button decreases magnification. In both cases, magnification changes continuously throughout the duration of the mouse press and is centered on the

location of the click. Alternatively, a slider in the GlassEye toolbar can be used to directly adjust the magnification.

Animation in visual interfaces can be useful for illustrating system state and for making interfaces more comprehensible [2]. In zoomable interfaces, animation has been shown to be potentially beneficial for tasks involving the reconstruction of tree structures based on exploration experience [3]. GlassEye uses animation to illustrate the transitions associated with change of magnification (or change of position through hyperlinks). In addition to providing users with orienting information that may assist navigation, these animations provide an engaging visual effect that may encourage exploration.

5.2 Navigation

The use of varying magnification in the GlassEye browser presents some challenges for effective navigation. As the magnification level increases, labels for the individual compositions grow larger, eventually disappearing once the magnification level exceeds their maximum magnification. In essence, the view "zooms past" the objects. This may lead to a

Figure 4: Detailed display for a specific composition.

view in which the user does not have any navigational cues – a situation known as “desert fog”[13].

The GlassEye browser uses an overview window to address this problem. This display contains a miniaturized version of the initial contents of the main GlassEye window. However, unlike the display in the main window, the overview does not change magnification: as the user zooms in or out in the main window, the scale of the overview remains constant. A field-of-view box [20] in the overview illustrates the position of the current main window view in the xy space of the original display: as the magnification increases, the size of the outline in the overview window decreases (and vice-versa). The overview can be used for navigation – clicking anywhere in the outline view will cause the view in the main window to move to a view centered on the position selected.

The overview window provides users with sufficient context to avoid the “lost in zooming space” problem. The combination of the miniaturized representation of the overall space and single-click navigation provides users with a tool that can always be used to find some area of potential interest. As screen space constraints limit the size of the overview display, items in the overview are generally not legible: they are simply intended as indicators of the presence of an item of potential interest at a given location. As such, they provide additional, if incomplete, navigational residue [11].

A “reset” button on the GlassEye tool bar provides an additional navigational aid, in the form of a single button that can be pressed to restore the browser to the original view.

A simple search facility provides support for finding compositions by name. When a user searches for targets that match a given string, matching labels are outlined and highlighted while all others are grayed out. This provides a display that supports quick visual location of the objects of interest while retaining appropriate context. Currently, search is limited to simple textual prefixes: a production system would be likely to support more sophisticated searching.

5.3 Alternate Views

Spotfire [25] supports starfield displays based on arbitrary axes: given an n -dimensional data set, starfield displays can be built using any two of those dimensions as the spatial axes. This feature provides significant flexibility for users interested in identifying potentially interesting visualizations. However, many of the possible views of a given data set are not useful, leaving users with the challenge of sorting through a potentially large space in search of appropriate views.

The GlassEye browser provides a small set of pre-defined views to provide the flexibility of multiple views of the data set without the potential for confusion associated with arbitrary choices. As mentioned above, the default display of the GlassEye browser uses the date of composition on the x -axis and the number of instruments on the y -axis. A list-box at the top of the screen

can be used to select one of the alternative views: currently, three choices are provided.

The current implementation of GlassEye uses the alternative views to explore a potentially interesting approach to visualization of a catalog of music: the use of subjective criteria to arrange items in a visualization. Given the potentially evocative nature of music, we hypothesized that some users might want to use criteria such as intensity, mood, or other less concrete descriptors in order to identify pieces of music.

Toward that end, we assigned ordinal values of “mood” and “intensity” to each of the compositions in the GlassEye browser. These attributes formed the basis for the second view of the data, which we labeled the “emotion finder”: mood on the x -axis and intensity on the y -axis. This might be used to identify compositions with a desired mix of these features: for example, “which pieces are upbeat but not too intense?”

A view displaying popularity (x -axis), vs. length (y -axis) is provided as a third alternative. Popularity values might be useful for individuals interested in locating pieces that might be most familiar.

5.4 Music

GlassEye uses a generalization of the Musical Scrollbar to support navigation through Glass’ music. Each composition has an associated audio sample, which is played when the user mouses over any component of the piece, including both the top-level label and the details presented at greater levels of magnification. A text label at the top of the GlassEye browser displays the name of the piece being played, providing users a means of identifying pieces of interest.

The use of mouse-over to initiate audio playback provides a highly interactive, if potentially confusing, model. Moving quickly throughout the system, users might hear only a few seconds of any single piece, but this might be sufficient for identifying items of potential interest.

Of course, incessant playback might be annoying in certain circumstances. Musical playback in GlassEye is turned off by default and must be activated through a checkbox selection.

6. DISCUSSION

The GlassEye browser synthesizes a variety of information visualization techniques for use in a novel domain. Where past starfield visualizations such as the HomeFinder [28] and the FilmFinder [1] have focused on largely data driven tasks, with known goals, GlassEye attempts to provide support for more open-ended exploration. Towards that end, animated zooming, color highlighting, and digitized audio are used to entice users. Our hope is that the combination of these features with the dynamic querying and spatial representation of the starfield view will help users build an understanding of the data while enjoying the experience of exploring it.

In many ways, this system might be most valuable as a test platform for ideas. Several of our design decisions arose out of intuition, and were therefore not justified by user studies or existing research. Empirical evaluation of these aspects of the design might prove interesting and constructive for deployment of future systems based on the GlassEye browser. Furthermore, the current design identified numerous issues that must be addressed in any eventual production versions of the browser.

6.1 Navigability

Although providing an enjoyable user experience is one of the goals of the Glass site, the GlassEye browser will be most successful if users are consistently able to locate items of interest. The starfield display, query filters, overview window, and hyperlinks combine to provide navigation support that is potentially both powerful and confusing. Our initial use of the browser identified several design issues that affect navigability: further examination of these issues might inform subsequent revisions to the system.

6.1.1 Representation and Layout of Data

As described above, the top-level representation for a composition is currently either the text of its title, or a circle if no title exists. The Glass catalog contains a large number of untitled works: inclusion of these works in the browser might lead to a display containing large numbers of circles. In this case, additional measures or representations might be needed to support navigation and searching.

The use of title strings to represent compositions raises additional concerns relating to space utilization and overlap. Currently, these strings are simply placed by calculating the appropriate (x, y) position of the composition in the currently selected axes, and then using that position as the starting point for drawing the text. This simple approach creates two potential problems. The first – interpretation of a text string occupying significant area in an xy plot – is likely to be a problem for any use of textual representations. For example, if the text of a label covers the horizontal space corresponding to the years 1973-1978 inclusive, what conclusions can the user draw about the date of composition?

The second, more serious problem involves label overlap. Even with the small subset of the Glass catalog used in the current prototype, many of the labels overlap in a manner that obscures text and inhibits comprehension. Although the extent of this problem might be lessened through algorithms that attempt to avoid overlap, such approaches would have the detrimental side effect of moving labels away from the x - y position implied by the values of the attributes used in a given view of the data. This would reduce the information conveyed by the position of the label in the plane. Labeling approaches such as excentric labels [9], label connectors, or semantic labeling [15] might provide alternative possibilities that would avoid this loss of information.

A final layout concern involves the display of detailed information for a given composition. In the current implementation, detailed data about each composition is displayed when the user zooms to an appropriate magnification. However, the layout of this information is somewhat informal, and might be improved through appropriate graphic design.

6.1.2 Searching

Search facilities in the browser are currently limited to simple text prefix matching based on titles of individual compositions. Extensions of this facility to include searches through detailed data and/or structured searches restricted to particular data fields would significantly increase the utility of the searching tools.

Such extensions would also bring additional design challenges relating to the display of search results in a multi-scale interface. Specifically, how should search results be displayed if a matching target is not visible at the current magnification? If the user can zoom in to see the target, a “critical zone” indicator [13] might be displayed to indicate the presence of material that would not be visible without further magnification. In the dual case – if the target is not visible because magnification is too great and the view has “zoomed past” it – appropriate strategies for indicating the search match are not obvious, and would be an interesting area of further investigation.

6.1.3 Aggregation

Musical works can often be viewed as aggregations of subcomponents: symphonies consist of movements, operas of arias, etc. Display of the hierarchical nature of many of the pieces in the Glass catalog in a manner that supports effective navigation raises many additional issues that are not addressed in the current prototype.

Magnification and zooming might be used to illustrate hierarchical relationships. From the beginning overview screen, users might zoom in to a display that shows each of the components of a given work, along with background description that pertains to that work as a whole. Further zooming might lead to detailed information about each of the individual components. This approach might be fruitful, but users may become disoriented if cues are not provided to help place the selected component within the greater whole.

Appropriate handling of hierarchical works is also a concern for the layout of the starfield display. The coordinate axes used by the GlassEye browser rely on the simplifying notion that a given work can be accurately categorized along two dimensions. However, for works that have multiple subcomponents that differ in length, instrumentation, mood, and other criteria that might be used on a starfield display axis, it is not at all clear how the work as a whole should be described.

One possibility might be to present the work as occupying some area in the 2D space, corresponding to the range of values that might be appropriate for the work in a given attribute.

This approach is likely to be effective only if the values are continuous, as discontinuous values might require multiple areas. Another possibility is the provision of an option that would allow display of individual sub-components as works in their own right. This promotion of the components of a given work would allow for identification of individual components, with the possible costs of increased screen clutter due to a vastly larger set of data. Furthermore, it is not clear how such a display would indicate which components belonged to the same work, thus potentially increasing user disorientation and confusion.

The goal of any such approach would be to maximize the amount of data displayed without hindering comprehension or navigation. As this is a general tradeoff faced by information visualization systems, further examination in the context of the Glass catalog browser may be an interesting area of investigation.

6.2 Subjective Data

Starfield displays generally use well-defined attributes for the primary x-y axes. Numeric, alphabetic, or other familiar values are preferred, as they likely to be most meaningful to users and thus easily interpreted. Following this model, the GlassEye prototype uses attributes such as year, number of instruments, and length for some of the starfield views. However, other views – such as the mood vs. intensity display – break from this approach, relying instead upon subjective data.

Our interest in presenting views based upon subjective attributes was motivated by our belief that certain users or classes of users might be interested in identifying music that meet certain emotional or otherwise hard to quantify criteria. Some listeners might be interested in identifying pieces that were particularly light and “up-beat”, while others might be looking for pieces that were relaxing and “not very intense”. These subjective attributes were seen as particularly relevant for users interested in identifying music for licensing. For example, filmmakers might use these attributes to identify compositions with desired mood or tempo.

Attributes such as these might be displayed using technical measures or established musical terminology. While these measurements might be useful for providing accurate information to musicians and other skilled users, they might prove confusing for others. Subjective ratings on artificial scales were chosen in the hopes of avoiding confusion that might be caused by overly technical terminology.

Further investigation will be needed to understand the issues related to the use of attributes that are not well defined or understood as the bases for starfield visualizations. For instance, such axes may prove useful for establishing ordinal relationships, but users might find the meaning of distance on a particular axis harder to comprehend. More plainly, on a time-based scale, most users- will quickly infer that an item is roughly 45 minutes long if it is midway between the 30-minute item and the 60-minute item. However, the same distance might be much harder to interpret if

the scale is “mood” or “intensity”. Determination of when – if at all – such attributes are useful in starfield visualizations would be an interesting area of further study.

From a purely pragmatic viewpoint, these subjective ratings present a challenge when working with a catalog that may contain over one thousand entries. Generating these ratings and insuring their consistency is likely to be a daunting task. Ideally, those most familiar with Glass’ music – his production staff - would perform this task. Alternatively, future versions of a Glass site might use a collaborative filtering approach [22] to support visualizations based on ratings submitted by site visitors.

6.3 Music

Despite the inclusion of audio clips, the GlassEye prototype is primarily a visual system. This is not surprising, given the influence of information visualization techniques on the system design and the overwhelmingly visual nature of most current web environments.

Ideally, the musical content would be an integral component of a browser for the Glass catalog. Additional audio might be added to the current system in the form of longer audio clips presented along with the details for individual pieces. Other possibilities include additional facilities that support navigation through the catalog in terms of musical similarity. For example, the Glass catalog might be arranged by trends in the composer’s career, separating early repetitive pieces from more recent melodic compositions.

As browser capabilities and formats for storage and manipulation of music evolve, additional facilities might support more active forms of exploring the Glass catalog. For example, appropriately designed interfaces might support display of musical notation and investigation of the effects of changes to a score – perhaps providing a “write your own Philip Glass music” tool.

Information visualization tools attempt to provide users with graphic tools for visual exploration data sets, in the hopes of building understanding of patterns and detail. In some sense, the GlassEye browser represents a primitive attempt at “audio visualization”: the tools and techniques required to build a mental model of a body of music.

6.4 Scalability

The prototype browser currently includes a small subset of the Glass catalog, with all relevant data hard-coded into the Java applet. Although a production system would clearly require a more scalable data representation – perhaps a relational database or XML implementation – the visualization aspect of the browser faces scalability concerns. As the data set grows, crowding and overlap in the visualization space will increase, leading to decreased utility and increased confusion. Segmenting the data into different data sets – perhaps by type of composition – may

lessen this problem, at the expense of eliminating the comprehensive overview of the complete catalog.

6.5 Performance

Response times of 100ms have been suggested as requirements for effective dynamic query systems [23]. Generalizing this requirement to the GlassEye browser, we have a goal of 100ms response times for zooming operations, initiation of musical playback, and other interactions. This may be a difficult goal to reach for an applet-based visualization. While performance of components such as Swing and Jazz [5] are constantly improving, zooming operations remain computationally expensive and potentially slow.

Retrieval of the audio clips required for browsing of the Glass audio may further hinder performance. The current prototype uses Sun .au files for audio samples. Future implementations may want to use more modern technologies such as MP3 or streaming audio to overcome the limitations of .au files. Decisions regarding audio formats should consider the bandwidth requirements of various digital audio formats, and their impact upon users with low speed connections. A tiered interface with differing bandwidth requirements might be provided to avoid excluding users who do not have fast Internet connections.

6.7 Next Steps

This prototype provides a demonstration of the potential of combining starfield displays with Zoomable User Interfaces. We were particularly encouraged by the interactive, animated zooming and “fly-over” musical playback, which work together to create an enjoyable user experience. Unfortunately, the prototype raises almost as many questions (described above) as it answers. Additional efforts in several areas would be necessary for a production system based on the GlassEye prototype:

Visual design and placement within an overall site design. The look-and-feel of the browser window and controls, and the detailed display view (Figure 4) would clearly benefit from collaboration with a visual designer. Furthermore, this browser is likely to be used as one component of a larger site: as such, it should have a visual design that is consistent with the rest of the site. Ideally, this visual design would help users identify the browser with the site, and to associate both with Philip Glass, thus maximizing the potential for “branding.”

Performance, and other technical issues. As described above, GlassEye’s extensive use of Jazz and Java raises significant concerns regarding deployment for general Web usage. Although appropriate design and implementation improvements might reduce some of this overhead, it is not at all clear that acceptable performance for users of medium and slow connections will be achievable. Audio output presents another challenge, as streaming audio or MP3 would almost certainly be preferable to low-quality .au files. Re-implementation in a browser plug-in environment such as Macromedia’s Flash [16] or equivalent might be one approach to increasing performance and audio quality.

Navigability and usability: Concerns raised above regarding navigability and scalability are only two examples of the general concern of usability. Although we have been encouraged by our use of the system, we have not conducted any formal or informal usability testing. Such testing might help us understand the usability of the Zoomable starfield space and identify limits of information density, in order to inform future redesign.

7. RELATED WORK

The chronological, starfield display used in the GlassEye browser has its roots in earlier projects such as the FilmFinder [1], LifeLines [19], and the National Digital Library Collection browser [21]. GlassEye differs from FilmFinder and LifeLines in its approach to the use of zooming for display of overview and detail information. These earlier systems use zooming to narrow focus in to a particular point of interest, with details presented in a separate window. GlassEye uses zooming to navigate in scale to the desired detailed information, with a separate, dedicated window providing the overview.

Kullberg’s Dynamic Timelines [14] used timelines and zooming in a 3D environment to present a history of photography. The GlassEye browser and Dynamic Timelines make similar use of zooming for access to increased levels of detail. Kullberg used a 3D display to provide additional contextual information that might enrich the user experience, at the possible expense of additional processing overhead that is likely to be prohibitive for web-based environments. Dynamic Timelines does not address the issue of overviews, leaving the possibility that users might become “lost in scale space”.

Fernström and Bannon’s BROWSE project [10] was the first to examine the possibility of starfield displays for browsing an abstract space of musical works. BROWSE uses Macromedia Authorware in conjunction with custom-built sound-drivers, to support exploration of the Fleischmann collection of Irish traditional music. BROWSE uses audio localization as a tool for navigation: up to eight objects close to the cursor can play simultaneously, with objects represented by sources arranged in an audio space according to their position relative to the cursor. This use of spatial playback of audio represents both intriguing possibilities and significant technical challenges for use in a web-based browser for Philip Glass compositions.

Perhaps closest to the goals of the Glass site, Sony Music’s licensing web site [24] uses Thinkmap [26] technology to present music that might be licensed in an associative net of concepts supporting search, query, and song samples.

Other recent efforts aimed at navigation and visualization of music include the ISEE3D musical sound browser [27], and Open Hypermedia for hypertext linking within audio streams [8]. Possibilities for navigating musical spaces include query by humming and other approaches to the use of audio input as search queries [6][17]. Finally, systems like the Brain Opera’s Singing Tree [18] and the WorldBeat baton [7] suggest innovative

approaches for the use of new input modalities for interaction with music-based systems.

8. CONCLUSIONS & FUTURE WORK

The GlassEye prototype browser uses a range of techniques including starfield displays, multi-scale zoomable interfaces, dynamic querying, and audio output to support the exploration of a catalog of musical compositions. This design represents a synthesis of experience with information visualization systems and analysis of requirements and possible users tasks for a Philip Glass web site.

The success of any future production use of a system based on this prototype will be dependent upon work that evaluates the appropriateness and efficacy of the interface design while resolving the technical issues of scalability and performance. Without appropriate user studies and further analysis of tasks and needs, it will be difficult to conclude that the design actually meets user needs.

In any case, the GlassEye browser does not comprehensively address systems goals and user tasks described above. Additional background data, tighter integration of Glass' music, structured search, and other interactive tools for exploring and understanding music would be necessary for fully meeting the stated goals. Of course, further understanding of users and their needs may motivate creation of distinct tools for different needs: although a single interface for all classes of users has a certain appeal, the result may be a "lowest common denominator" that fails to meet the needs of any user population.

Continuing improvements in digital music formats and available bandwidth will undoubtedly lead to increased possibilities for web-based interfaces to musical content. Current interfaces tend to treat musical content as a "second-class" citizen, focusing instead on the more tractable problems of visual display. Extensions of the GlassEye prototype to include novel mechanisms for interacting with and creating music [7][18], or exploring the mathematical structure of Glass' compositions [29] present intriguing possibilities that might help old fans and new listeners learn about and appreciate the music of Philip Glass.

9. ACKNOWLEDGMENTS

This work was done at IBM T.J. Watson Research Center, and follows from work done at University of Maryland under the support of IBM's University Partnership Program. Thanks to Mark Podlaseck for making this project possible, to Edith Schonberg, John Karat, Kurt Munkacsi, Don Christensen, and Jim Keller, for their support and assistance, and Ben Bederson, Ben Shneiderman, and Judith Yanowitz for their comments on early versions of this paper.

10. REFERENCES

[1] Ahlberg C., and B. Shneiderman, "Visual Information Seeking: Tight Coupling of Dynamic Query Filters With

Starfield Displays," *Proceedings of ACM CHI '92 Conference*, 313-317 (May 1992).

- [2] Baecker, R., I Small, I, and R. Mander, "Animation at the Interface," In B. Laurel, ed. *The Art of Human-Computer Interface Design*, 251-267. Addison-Wesley, San Francisco, CA (1990).
- [3] Bederson, B.B., and A. Boltman, "Does Animation Help Users Build Mental Maps of Spatial Information?" University of Maryland CS-TR-3964, UMIACS-TR-98-73. <ftp://ftp.cs.umd.edu/pub/hcil/Reports-Abstracts-Bibliography/98-11html/98-11.html> (1998).
- [4] Bederson B.B., J.D. Hollan, K. Perlin, J. Meyer, D. Bacon, and G. Furnas, "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics," *Journal of Visual Languages and Computing*, 7, 3-31 (1996).
- [5] Bederson, B., and B. McAlister, "Jazz: An Extensible 2D+Zooming Graphics Toolkit in Java," University of Maryland CS-TR-4015, UMIACS-TR-99-24. URL: <ftp://ftp.cs.umd.edu/pub/hcil/Reports-Abstracts-Bibliography/99-07html/99-07.html> (1999).
- [6] Blackburn, S. and D. DeRoure, "A Tool for Content-Based Navigation of Music," *Proceedings of ACM Multimedia '98 Conference*, 361-368 (September, 1998).
- [7] Borchers, J., "World Beat: Designing a Baton-Based Interface for an Interactive Music Exhibit," *Proceedings of ACM CHI '97 Conference*, 131-138 (March, 1997).
- [8] DeRoure, D., S. Blackburn, L. Oades, J. Read, and N. Ridgway. "Applying Open Hypermedia to Audio," *Proceedings of ACM Hypertext '98 Conference*, 285-286 (June, 1998).
- [9] Fekete, J.D. and C. Plaisant, "Excentric Labeling: Dynamic Neighborhood Labeling for Data Visualization," *Proceedings of ACM CHI '99 Conference*, 512-519 (May, 1999).
- [10] J.M. Fernström, J. M., and L.J. Bannon,, "Explorations in Sonic Browsing," *Proceedings of BCS HCI '97, Bristol, UK*. <http://www.ul.ie/%7Eidc/library/papersreports/MikaelFerstrom/hciuk97/sonicbrowse.html>.
- [11] Furnas, G., "Effective View Navigation," *Proceedings of ACM CHI '97 Conference*, 367-374 (March, 1997).
- [12] Hochheiser, H., and B. Shneiderman, "Performance Benefits of Simultaneous over Sequential Menus as Task Complexity Increases," *International Journal of Human Computer Interaction*, forthcoming.
- [13] Jul, S., and G. Furnas, "Critical Zones in Desert Fog: Aids to Multiscale Navigation," *Proceedings of ACM UIST '98 Conference*, 97-108, (May, 1998).
- [14] Kullberg, R., "Dynamic Timelines: Visualizing the History of Photography," *Proceedings of ACM CHI '96 Conference Companion*, 386-387 (April, 1996).

- [15] Li, J., C. Plaisant, C., and B. Shneiderman B., "Data Object and Label Placement for Information Abundant Visualizations," *Proceedings ACM Conference on New Paradigms for Information Visualization*, 41-48 (November, 1998).
- [16] Macromedia Flash
<http://www.macromedia.com/software/flash/>
- [17] McNab, R., L. Smith, I. Witten, C. Henderson, and S. Cunningham, "Towards the Digital Music Library: Tune Retrieval from Acoustic Input," *Proceedings of ACM Digital Libraries '96 Conference*, 11-18 (March, 1996).
- [18] Oliver, W., J. Yu, and E. Metois, "The Singing Tree: Design of an Interactive Musical Interface," *Proceedings of ACM DIS '97 Conference*, 261-264 (August, 1997).
- [19] Plaisant, C., B. Milash, A. Rose, S. Widoff, and B. Shneiderman, "LifeLines: Visualizing Personal Histories," *Proceedings of ACM CHI '96 Conference*, 221-227 (April, 1996).
- [20] Plaisant, C., D. Carr, and B. Shneiderman, "Image Browsers: Taxonomy, Guidelines, and Informal Specifications," *IEEE Software* 11(1), 33-52 (March, 1995).
- [21] Plaisant, C., G. Marchionini, T. Bruns, A. Komlodi, and L. Campbell, "Bringing Treasures to the Surface: Iterative Design for the Library of Congress National Digital Library," *Proceedings of ACM CHI '97 Conference*, 518-525 (March, 1997).
- [22] Resnick, P. and H. Varian, "Recommender Systems," *Communications of the ACM*, 40(3), 56-58 (March 1997).
- [23] B. Shneiderman, B., "Dynamic Queries for Visual Information Seeking," *IEEE Software* 11(6), 70-77 (November, 1994).
- [24] Sony Music Licensing
<http://www.sonymusic.com/licensing>.
- [25] Spotfire <http://www.spotfire.com>
- [26] Thinkmap
<http://www.plumbdesign.com/projects/thinkmap.html>
- [27] Vertegall, R. and B. Eaglestone, "Looking for Sound? Selling Perceptual Space in Hierarchically Nested Boxes," *Proceedings of ACM CHI '98 Conference*, 295-296 (April, 1998).
- [28] Williamson, C., and B. Shneiderman, "The Dynamic HomeFinder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System," *Proceedings of ACM SIGIR '92 Conference*, 338-346 (June, 1992).
- [29] York, W. "Form and process," *Sonus* 1/2 (1981) reprinted in Kostelantz, R, ed. *Writings on Glass*. University of California Press, Berkeley CA (1997).