

A Taxonomy of Multiple Window Coordinations

Chris North and Ben Shneiderman*

Human-Computer Interaction Laboratory, Institute for Advanced Computer Studies,

*Institute for Systems Research & Department of Computer Science,

University of Maryland, College Park, MD 20742 USA

north@cs.umd.edu, ben@cs.umd.edu

<http://www.cs.umd.edu/projects/hcil/>

ABSTRACT

In current windowing environments, individual windows are treated independently, making it difficult for users to coordinate information across multiple windows. While coordinated multi-window strategies are increasingly used in visualization and web user interfaces, designs are inflexible and haphazard. The space of such linked-window strategies is not well understood and largely unexplored. This paper presents a taxonomy of coordinations, identifies important components, and reviews example interfaces. This 2x3 taxonomy provides guidelines for designers of applications, user interface toolkits, and window managers. We hope to encourage construction of generalized, end-user programmable, robust, multiple-window coordination capabilities.

KEYWORDS

User Interface, Coordination, Taxonomy, Multiple Window Strategies, Information Visualization, Window Management

INTRODUCTION

Users are dealing with increasing quantity, variety, and interrelated-ness of information. User tasks are becoming more complex, simultaneously drawing on many information sources. At the same time they are attempting more ambitious projects that demand rapid access to large amounts of visual information in many windows. Unfortunately, however, window managers and graphical user interfaces have not kept up with the users' needs and possibilities for visualizing information. In current windowing environments, windows are treated as independent and isolated, leaving users with the tedious job of individually manipulating one window at a time, even when separate windows are related by content or task. Hence, while interfaces have advanced from modal to many-window displays, these displays are still window modal in terms of interaction. Because of this lack of innovation, advantages of large screens and fast displays are lost or misused, leading to confusion, poor user performance, frustration, and missed opportunities.

In designing user interfaces for information exploration, a strategy that has proven increasingly effective is *multiple window coordination* [Shn98]. In this strategy two or more separate windows containing related information are used cooperatively for the accomplishment of a single task. The windows respond in a unified manner to user activity in the task domain. This strategy breaks the traditional

independent windows paradigm by *tightly coupling* interface components such that user interaction in one window has direct visual affects in other coupled windows, extending the direct manipulation approach across the once impenetrable window borders. Research indicates that potential gains from coordination strategies are significant, dramatically improving user performance (browsing software [SSS86], large 2D spaces [BW90], hierarchical tables of contents [CWM94], managing windows [KS97]).

In addition to specialized visualization needs, common tasks in multiple-window environments often practically beg for coordination. For example, consider the task of authoring a large web document (see Figure 1). Scrolling the web page window might produce a synchronized scroll of the corresponding HTML source view to show the same portion of the document, providing simultaneous views of WYSIWYG layout and detailed layout description. For quick access, selecting a chapter title in a separate table-of-contents view could scroll both of the previous views to that chapter's text. Conversely, scrolling the document views could highlight the corresponding section title in the table-of-contents view to help maintain context. Using the document as an index, selecting hypertext links could open the destination pages in a fourth window (e.g. to examine references in a technical paper).

This paper proposes a generalized multiple-window-coordination capability, in which users can establish coordination links between windows to create custom visual information-seeking environments. The link has 3 conceptual components: the 2 endpoints are the windows (or contained information) being coordinated together, and the connection represents a tight coupling between them. With several different types of coordinations, users can link groups of windows to operate according to the custom needs of the current task as in the web authoring example above. Users have the flexibility, for example, to add another document focus view which is also linked to the table of contents and destination page windows and to save this composite coordinated window group for future reuse with other documents or similar tasks.

To build such a system, a thorough understanding of the actions and objects involved and a language to describe such coordinations are needed. The following taxonomy identifies the major classes of multiple window coordinations and the tight couplings, components, and

interaction that support them. It also reviews examples of each with research systems from the literature as well as commercial systems. We focus on coordinations for visual information seeking tasks and attempt to show application to each of the seven basic datatypes (1D, 2D, 3D, Mutli-D, temporal, hierarchic, network) [Shn96]. We also demonstrate these with mockups for browsing text documents (as in the above web example), and 2D and 3D data (as in the Visible Human Explorer [NSP96]). See [NWS86], [SSS86], [CPF84] for taxonomies of general multiple window strategies, and [CR96], [Twe97] for taxonomies of interaction in information visualization.

TAXONOMY OF MULTIPLE WINDOW COORDINATIONS

The taxonomy (Table 1) organizes coordination strategies into a 2x3 matrix. First, we define a window to consist of two major components, a *view* and a collection of information *items*, each of which has an associated basic user action:

- *Selecting items*: Windows contain collections of information items of any of the seven basic datatypes (e.g. characters, words, paragraphs of 1D text, or pixels, regions of a 2D image, or nodes of a directory tree). In direct manipulation, the basic user action on items is selection.
- *Navigating views*: A window's view provides a visualization of the collection of items. Since the view might show only a small portion of large collections, users navigate the view to see other parts. Different navigation actions apply to different data types (e.g. scroll, pan, zoom, slice, rotate, ascend/descend tree, follow link, open file, etc.).

Coordinating a pair of windows tightly couples one of these actions in the first window to another in the second window. The primary dimension of the taxonomy classifies coordinations by the three possible combinations:

1. Selecting items ↔ selecting items
2. Navigating views ↔ navigating views
3. Selecting items ↔ navigating views

The second dimension classifies by whether the collections of information items contained in the coordinated window pair are the same or different:

- A. Same collection of information items in each window:** Coordinating multiple representations (windows) of a single collection (e.g. a bar chart and tabular list of the same set of data points) allows users to simultaneously browse several different aspects of the information. In this case, there is an implicit identity relation between the corresponding items between each view.
- B. Different collections of information items in each window:** The collections are not identical but are interrelated (e.g. a document and a dictionary). Relationships must be explicitly defined either in the information itself, by a separate database or matching algorithm, or by the user.

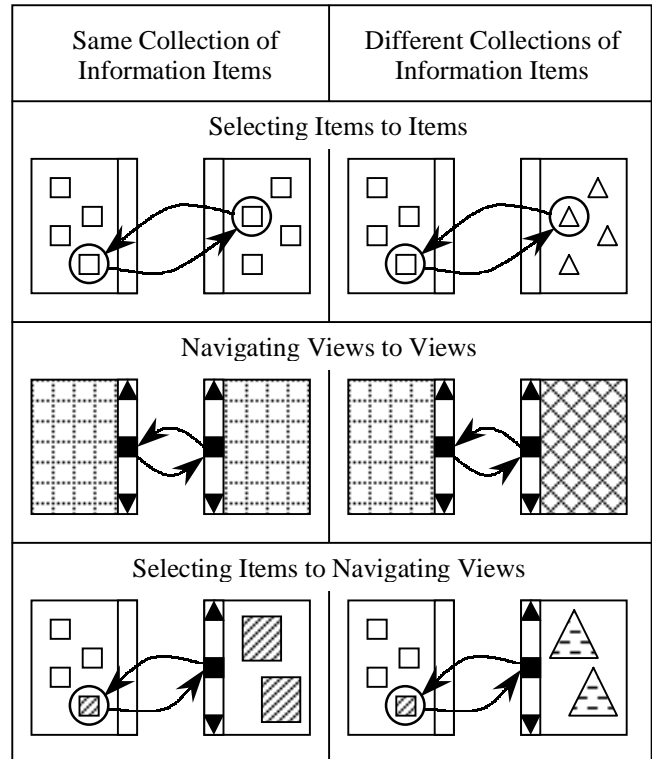


Table 1. A 2x3 taxonomy of multiple window coordinations.

1. Selecting Items to Items

This coordination tightly couples selecting items in one window to selecting items in another window, to help users correlate equivalent or related items.

A. Same information collections

When users select (highlight, paint, “brush”) an item (or set of items) in one view, the system immediately highlights the equivalent item (or set), representing the same underlying data elements, in the other views. Many data analysis systems (e.g. Spotfire [IVE97], XGobi [BCS96]) use this coordination to visualize high-dimensional data point sets with multiple coordinated 2D scatter plots ([MSB90] credits invention to C. Newton in 1978). These systems provide many options for selection method (single, cumulative, box, lasso, etc.) [Wil96] and highlight style (color, size, glyph, etc.). Visage [RLS96] generalizes this coordination, allowing users to drag-and-drop and then brush tabular data elements among many types of visualizations, including charts and geographical maps. The Navigational View Builder [MFH95] brushes nodes in hierarchical information, linking treemaps (emphasizing numerical and categorical attributes), ConeTrees (emphasizing structure), and outliners (emphasizing node names). With Lilac [Bro91], a two-view document editor, selecting text in the WYSIWYG page view also selects the corresponding text in the source text view, as in our web example (see Figure 1).

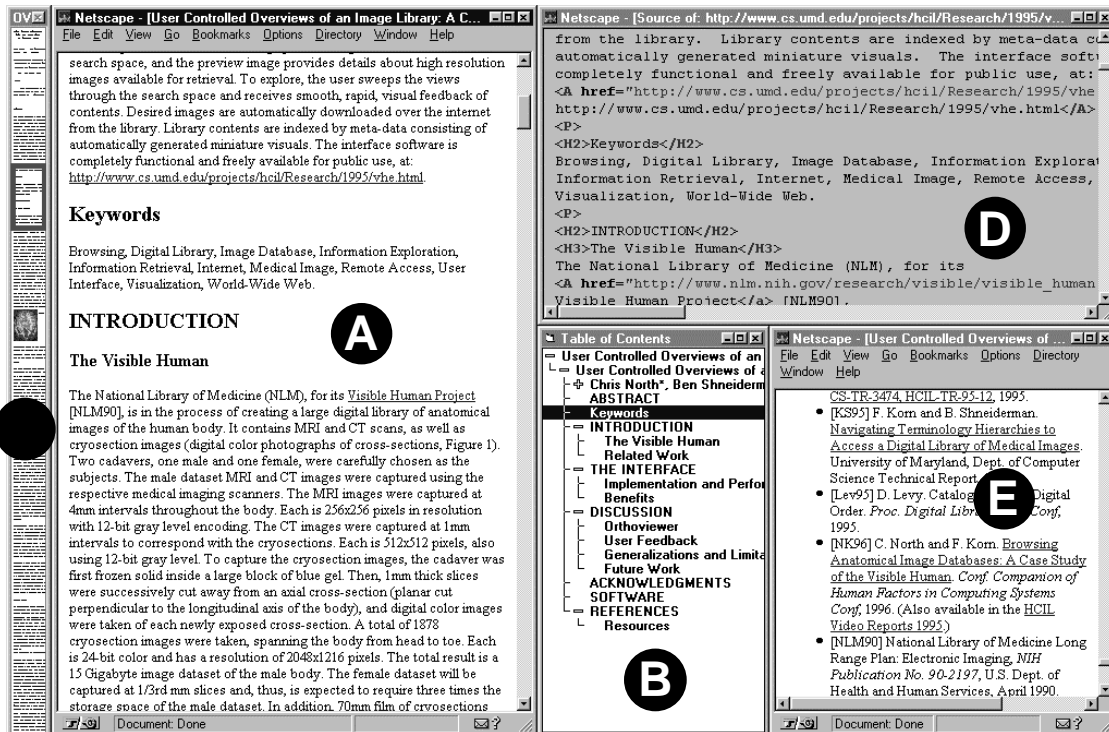


Figure 1. Coordinated windows for browsing web documents: (A) the main page window, (B) a corresponding hierarchical table-of-contents window with item selection tightly coupled to main page view, (C) a corresponding zoom-out page overview with a field-of-view box tightly coupled to the main page view, (D) a synchronized-scrolling HTML source view, and (E) an additional page view that opens links selected in the main view.

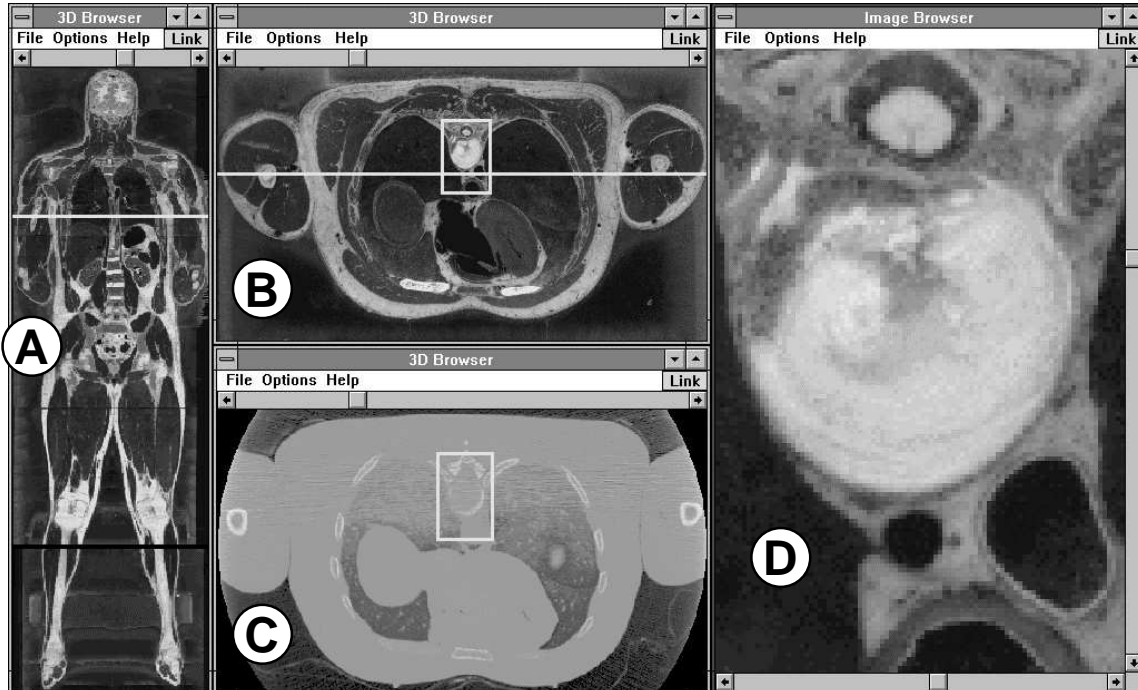


Figure 2. Composition of several coordinations for browsing the Visible Human 3D image data. Windows show: (A) vertical cross-section, (B) horizontal cross-section, (C) CT scan, and (D) horizontal detail cross-section of the spine. Coordinations include: Intersection lines in (A) and (B) navigate the views, (B) and (C) navigate synchronously, and field-of-view in (B) manipulates (D).

This strategy also employs more advanced forms of selection. Filtering, instead of highlighting selected items, removes unselected items from display (elision). Dynamic Queries [AS94] allows users to quickly select subsets of data based on attributes using widgets such as range-sliders, and filter to the subsets in other coordinated visualizations. With the LinkWinds system [JBO94], users can link control windows to visualization windows for dynamic filtering of scientific sensor data. The Attribute Explorer [Twe97] uses additive encoding of multiple attribute selections to highlight each item in a different color depending on the number of the selections it satisfies.

B. Different information collections

Brushing can be generalized to visualize not just equivalencies between same information collections, but also arbitrary relations between different collections. For example, interrelated data records could be brushed by consulting a relational database during the brushing process (Figure 4). In Microsoft Word, selecting an annotation in the annotations window also highlights its corresponding pointer in the document text window to help users locate it in the text. In extending the Visible Human Explorer [NSP96], users could brush a view of the cross-section image data to view a list of corresponding textual annotations in the MeSH (Medical Subject Headings) controlled vocabulary hierarchy (Figure 3).

2. Navigating Views to Views

This coordination tightly couples view navigation in one window to simultaneous view navigation in another window. This maintains synchronization of views while navigating (e.g. panning, zooming, slicing, traversing, etc.) multiple views (generally of essentially equivalent levels of detail) through correlated large information spaces. A scroll bar is a typical window navigation mechanism (panning).

A. Same information collections

Synchronized scrolling tightly couples the scroll bars of 2 windows. For example, in the case of same information collections, when editing a document (1D) using a markup language (e.g. HTML), it is helpful to simultaneously view and edit both the markup text representation (for fine-grain control, macros, etc.) as well as the WYSIWYG (e.g. web page) layout representation (to see the resulting output). Automatically scrolling both windows, when the user scrolls either, to synchronously navigate the space, avoids losing the relationship between representations as well as tedious repetition of navigation actions in each window by the user (see Figure 1 A,D). WordPerfect uses this strategy to display formatting codes in a separate window.

B. Different information collections

In the case of different information collections, Microsoft Word provides synchronized scrolling with document, endnote, and annotation windows. Scrolling the document causes the endnote and annotation lists to scroll to show those that are referenced within the current document view. Although if the number of items currently in view in the annotation window does not equal the number referenced

within the current view of the document window, it is not clear whether it favors those on the top or bottom of the views. With Logos Bible Software, users can link sets of windows to simultaneously scroll through several Bible translations, commentaries, and study guides, which all share a common ordered hierarchical structure of book, chapter, and verse.

DEVise [LRB97] generalizes the Navigating Views to Views strategy to 2D, allowing users to link the x and y axis between multiple graphs to show the same domains and ranges. This coordination is easily extendable to 3D, panning correlated cross-section, CT, and MRI images through the human body (Figure 2 B,C) (see also [CSP97]).

3. Selecting Items to Navigating Views

This coordination tightly couples item selection in one window to view navigation in another window, and vice versa. Users can select items in context to locate and view corresponding detailed information in separate windows.

A. Same information collections

Hierarchical Browsing uses same information collections. Overviews provide a global map of the whole collection, and detail views provide detailed information about a small portion of the collection. Coordinating the views indicates the location of and provides a mechanism for navigating the detail view from within the context of the overview. This is advantageous over detail-only browsing since overviews indicate what information is available, provide context for details, guide browsing, promote exploration, and help avoid getting lost. This strategy contrasts with distortion oriented techniques, which attempt to show details within the context of the overview in a single window by distorting the view. An important metric is the zoom factor between the overview selection and detail view. Larger zoom factors provide for more information. While zoom factors for distortion techniques are typically limited to 5 or less, coordinated windows can reach zoom factors of 20 [PCH92]. Also, several of these coordinations can be chained together using intermediate views [PCS95] to reach zoom factors of 400 and so on.

This coordination can be further segmented by the continuity of the information space. For *continuous* spaces, dragging or resizing a *field-of-view* indicator (selection) in the overview is tightly coupled to pan or zoom (navigation) the detail view (Figure 2 B,D), and vice versa (bi-directional coupling). Scroll bars, albeit poor overviews of their associated main window, are a simple 1D example. The Information Mural [JS95] and SeeSoft [BE96] provide highly reduced images of large documents or software code, using color coding and anti-aliasing algorithms, and a proportional field-of-view instead of the standard “thumb” control, for navigating document windows (similar to Figure 2A). The “cursor” link in DEVise [LRB97] links a 2D field-of-view in an overview plot to the panning control of the axes in a detail plot. Plaisant et al. [PCS95] develops a formal notation for specifying this coordination for browsing large images in a 2D continuous space. For a 3D

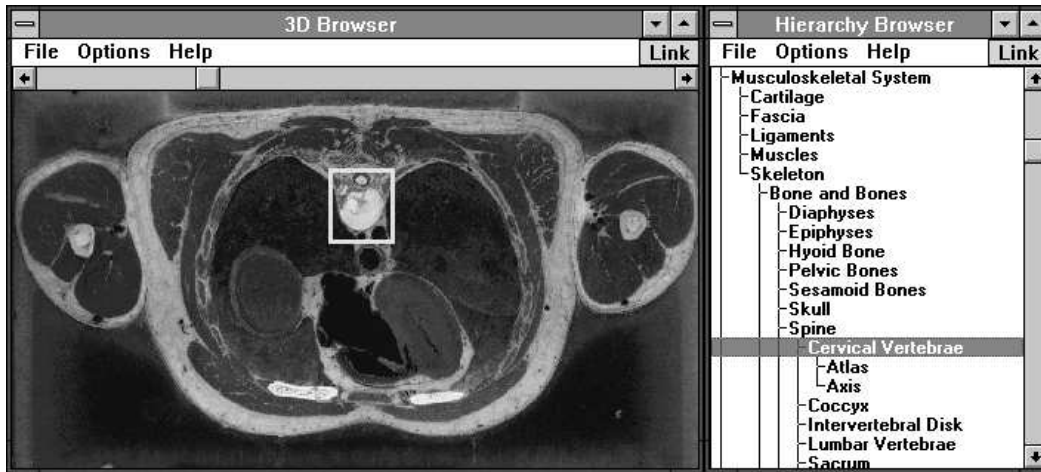


Figure 3. Coordination between Visible Human images and a textual hierarchical anatomy dictionary (MeSH) for brushing interrelationships. A field-of-view box in the image window highlights the structure corresponding to the selected term in the text window.

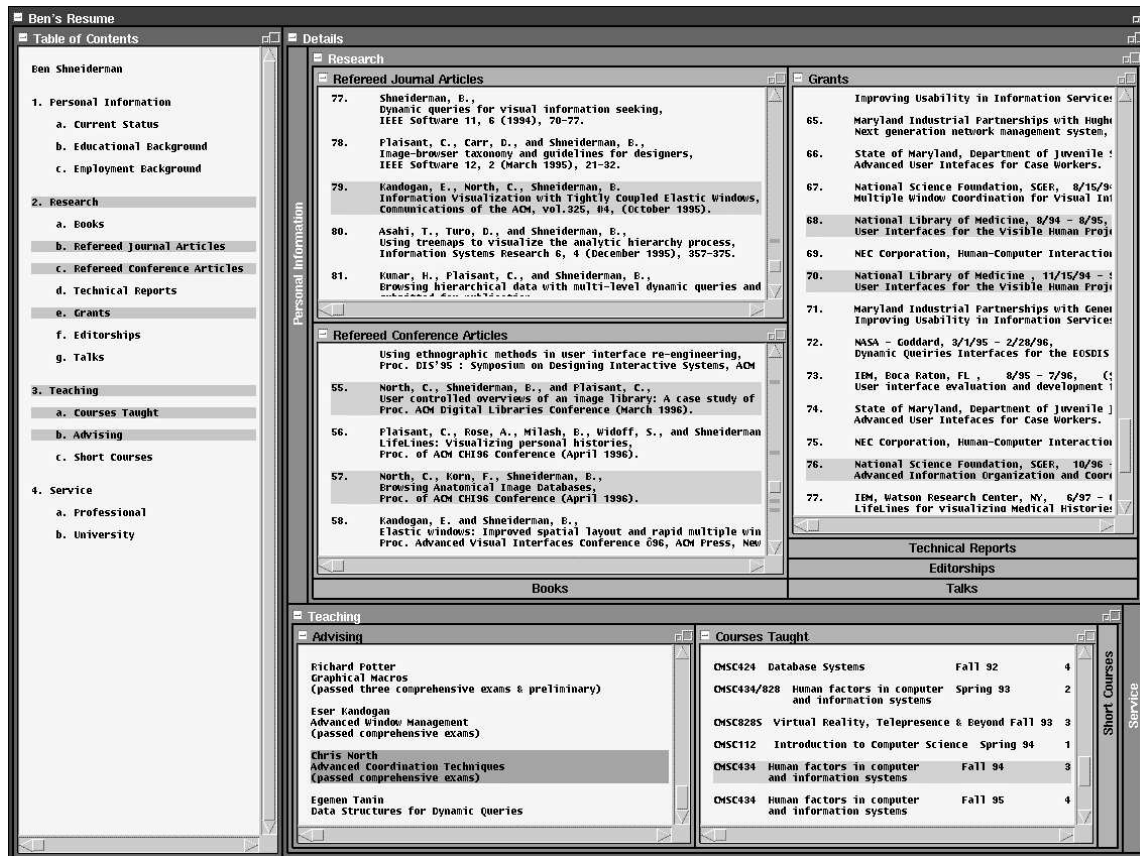


Figure 4. Coordinations for browsing a professional history along with Elastic Windows [KS97] hierarchical window management. Selecting items in the Table of Contents window navigates the window hierarchy, opening and closing corresponding nodes. Users can view relations between items in these windows. For example, when a student's name is selected, associated papers and grants are highlighted. Off-screen highlights are indicated in the scroll bars of windows.

volumetric image space, with the Visible Human Explorer [NSP96] users can rapidly navigate each orthogonal 2D cross-section view through the human body by dragging the corresponding cut lines in the other views, and receive continuous feedback of contents (as in Figure 2 A,B).

Discrete spaces, such as hierarchical menus, use *highlighted selection*. Users select individual items, which are subsequently indicated by highlighting, in an overview window (e.g. table of contents of a book, index web page for department faculty, computer directory listing) to navigate to corresponding details in an adjacent window (e.g. page of the book, professor's web page, contents of a file, respectively) as in Figure 1 A,B. In a real time environment, users can quickly flip through many pages by simply dragging the mouse over a series of items in the overview. With web site visualization tools, such as the Navigational View Builder [MFH95], users can select any node on a large site to display that web page in a separate browser window. This strategy is especially popular in pane-based systems (SpotFire [IVE97], Windows Explorer, Netscape frames, etc.). A variant of this approach shows details of selections in a new popup window instead of a given static window, as in the FilmFinder [AS94] (selecting a dot on a scatter plot displays that record's fields, including pictures). In the Apple Dylan programming environment [DP95], users can link multiple instances of its code browser to navigate several levels of object-oriented constructs.

B. Different information collections

In the case of different information collections, directly selecting an item in one context shows additional details or related information from other sources in another window. This case focuses more on bringing together many sources of information rather than selecting from overviews, as in web searches, a book index, or technical report references. For example, while a technical paper is not necessarily an overview of its references, directly selecting a reference within that context might show the linked document in another window (Figure 1 A,E). Similarly, Microsoft Word can look up selected words in a dictionary, thesaurus, or other reference material. More advanced capabilities are needed. For example, selecting a phrase on a web page could automatically initiate a web search, perhaps offering a choice of specific resources to search, such as yellow pages, maps, or medical databases.

ADVANCED COORDINATION

Some systems allow users to interactively link a set of windows within the system so users can choose which windows participate in specific coordinations. LinkWinds [JBO94] and Dylan [DP95] use a drag-and-drop action from source window to destination window to establish the link. DEVise [LRB97], XGobi [BCS96], and Logos use selection from lists. Those that implement multiple coordination options then use a dialog box to choose among the link types. LinkWinds draws lines between linked windows to visualize the links.

Composite coordinations can be built by chaining several coordinations together. Then, by the transitivity of tight couplings, user actions in one window propagate through the chained coordinations resulting in visual feedback in many windows. With this technique, designers can build very powerful comprehensive browsing environments, which may reveal interesting relationships in information. For example, in the Visible Human, dragging a cut line in one cross-section view would pan the corresponding orthogonal view, which, in turn, might pan a correlated CT scan view, which may scroll a high-resolution detail view, and update its field-of-view back in the original cross-section view (see Figure 2).

The coordinations described in this taxonomy could be extended for advanced data mining capabilities by hooking arbitrarily complex algorithms onto specific tight couplings. For example, in selecting items to items, when a user selects an item in one window, the resulting items selected in a coordinated window could be determined by a custom program. Examples include tools for visualizing file differences by highlighting the differing parts of two similar documents, and spell-checker dialog windows which iterate over the text of another window highlighting misspelled words and offering some choices.

While we have focused on multiple-window coordinations for visual information seeking tasks, there are other categories such as manipulation and editing tasks (consistency, coordinated layouts), dynamic data (automatic triggers, alarms), and collaborative multi-user environments.

IMPLICATIONS

The purpose of this taxonomy is to provide guidance for three major communities:

1. Application Designers:

First is the provision of general guidelines for user interface designers to: (a) identify opportunities to utilize coordinations in application designs to improve user performance and productivity, and (b) build them into systems (or rebuild faulty designs) in a robust and consistent fashion. For example, Netscape frames were initially well-received because users strongly desired hierarchical-browsing coordination capabilities in web site navigation. However, typical poor designs lack interactive robustness, leaving out important coordination elements, and lead to confusion. Common failures are: not highlighting current selection in overview frame (missing visual element), overview highlight not responding to navigation in detail view (bi-directionality failure), and unclear relation between frames and history keeping (lacks history of coordination). Similarly, Microsoft Word's coordination strategies for annotations and endnotes are inconsistent (for example, only annotations supports brushing). Application designers have done well to create standards for dialog boxes to ensure consistent look and feel, and should do the same for coordinations.

2. User Interface Toolkit Designers:

Typically, the reason application designers fail to build robust and consistent coordinations, if any, into systems' user interfaces is because of the lack of support in development toolkits. Toolkit designers have excelled at providing developers with ready-made widgets, automatic layout capabilities, and drag-and-drop construction as in Visual Basic. However, toolkits are still based on the independent-windows paradigm, lacking both the visual elements and the underlying support for tight coupling required for coordination. Programming each coordination is difficult, code intensive, prone to mistakes as previously described, and is a significant software engineering task. The visual elements, such as highlighting techniques, multiple selections, coordination affordances, and link visualizations should become standardized in (and easily invoke-able from) common toolkits. Data structures to represent the links and algorithms to support propagation of interaction across chained many-window coordinations can become elaborate and are also in need of standardization.

3. Window Manager / Operating System Designers:

The taxonomy suggests the viability of a tool for a general multiple-window-coordination capability, in which end users can establish coordinations, on the fly, by interactively linking windows and specifying coordination types and details. Such interaction across window borders may be the next big step in escaping the limitations of the current independent windows approach. The features of such a tool are:

- Integration into window managers, modeled after existing cut-and-paste and drag-and-drop coordination facilities, with simple window applications hooks, to provide *universal support* among all windows.
- Coordination within as well as between many *different data types*.
- User interface of coordinations: *Robust* and *consistent* look and feel of coordination across domains.
- User interface for coordinating (linking) windows via direct manipulation: a language or *notation*, based on the continued elaboration of this taxonomy, to provide a *specification* mechanism for coordination as well as *visualization* of coordinations in place. Possible approaches are textual spread-sheet style or visual languages (in the spirit of [PCS95] for specifying multi-window image browsers, [Bor86] constraint-based simulation components, [HM90] user interface component layouts, [ISB95] hypertext layouts). This gives users the *flexibility* to customize browsing environment to tasks.
- Capability to build up *composite* coordinations, by *chaining* many together, which can be saved or shared.

CONCLUSION

The benefits of multiple window coordination appear to be significant:

- *User performance improvement* by reducing information *access* time, especially in the Hierarchical browsing case (as in [SSS86], [BW90], [CWM94]), and by reducing user actions (*housekeeping*) required to maintain a synchronized state by hand (as in [KS97]). For example, with a simple key-stroke level analysis, scrolling correlated cross-section and CT scan images with synchronized scrolling coordination between the two views is half the work than without, and is even better during exploratory tasks in which users roam through the space extensively with dragging actions.
- *Discovery* of unforeseen relationships within information that otherwise would be difficult to find or completely overlooked without coordination. Helps improve users' *understanding* of the information by supplementing their mental model with the additional dimension of coordination.
- *Unification* of the desktop: We believe that it is often wiser to present complex information, not in a single integrated complex display, but in multiple simpler displays using coordination to navigate, maintain relationships, etc. (Figure 3). Is this then anti-Tufte, who often advocates 'integrated arguments'? Perhaps not; for coordination is *integration by interaction* rather than simply integration by visual design.

This taxonomy helps identify potential coordinations and provides guidance for three communities of designers. There are rich possibilities in

1. Selecting Items to Items
2. Navigating Views to Views
3. Selecting Items to Navigating Views

from the same and different collections of information items. Other coordinations may be possible, but this initial taxonomy can generate many user interface improvements to existing systems.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Grant No. NSF IRI 96-15534.

REFERENCES

- [AS94] Ahlberg, C., Shneiderman, B., "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays", *Proc. ACM CHI'94 Conference*, pp. 313-317, (1994).
- [BE96] Ball, T., Eick, S., "Software visualization in the large", *IEEE Computer*, 29(4):33-43, (April 1996).
- [BW90] Beard, D., Walker, J., "Navigational techniques to improve the display of large two-dimensional spaces", *Behaviour & Information Technology*, 9(6), pp. 451-466, (1990).
- [Bor86] Borning, A., "Constraint-Based Tools for Building User Interfaces", *ACM Transactions on Graphics*, 5(4), pp. 345-374, (October 1986).

- [Bro91] Brooks, K., "Lilac: A Two-View Document Editor", *IEEE Computer*, 24(6), pp. 7-19, (June 1991).
- [BCS96] Buja, A., Cook, D., Swayne, D., "Interactive High-Dimensional Data Visualization", *Journal of Computational and Graphical Statistics*, 5(1), pp. 78-99, (1996).
- [CPF84] Card, S., Pavel, M., Farrell, J., "Window-based computer dialogues", *Proc. INTERACT '84*, London, UK, pp. 355-359, (1984).
- [CSP97] Carlis, J., Safonov, A., Perrin, D., Konstan, J., "The Neighborhood Viewer: A Paradigm for Exploring Image Databases", *Proc. Companion of ACM CHI'97 Conference*, pp. 299-300, (1997).
- [CWM94] Chimera, R., Wolman, K., Mark, S., Shneiderman B., "An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents", *ACM Transactions on Information Systems*, 12(4), pp. 383-406, (Oct. 94).
- [CR96] Chuah, M., Roth, S., "On the Semantics of Interactive Visualizations", *Proc. Information Visualization '96 Conference*, IEEE, pp. 29-36, (1996).
- [DP95] Dumas, J., Parsons, P., "Discovering the Way Programmers Think About New Programming Environments", *Communications of the ACM*, 38(6), pp. 45-56, (June 1995).
- [HM90] Hudson, S., Mohamed, S., "Interactive Specification of Flexible User Interface Displays", *ACM Transactions on Information Systems*, 8(3), pp. 269-288, (July 1990).
- [ISB95] Isakowitz, T., Stohr, E., and Balasubramanian, P., "RMM: A Methodology for Structured Hypermedia Design", *Communications of the ACM*, 38(8), pp. 34-44, (August 1995).
- [IVE97] IVEE, Spotfire, <http://www.ivee.com/>.
- [JBO94] Jacobson, A., Berkin, A., Orton, M., "LinkWinds: Interactive Scientific Data Analysis and Visualization", *Communications of the ACM*, 37(4), pp. 43-52, (April 1994).
- [JS95] Jerding, D., Stasko, J., "The Information Mural: A Technique for Displaying and Navigating Large Information Spaces", *Proc. IEEE Symposium on Information Visualization*, pp. 43-50, (October 1995).
- [KS97] Kandogan, E., Shneiderman, B., "Elastic Windows: Evaluation of Multi-Window Operations", *Proc. ACM CHI'97*, pp. 250-257, (March 1997).
- [LRB97] Livny, M., Ramakrishnan, R., Beyer, K., Chen, G., Donjerkovic, D., Lawande, S., Myllymaki, J., Wenger, K., "DEVise: Integrated Querying and Visual Exploration of Large Datasets", *Proc. ACM SIGMOD'97*, pp. 301-312, (1997).
- [MSB90] McDonald, J., Stuetzle, W., Buja, A., "Painting Multiple Views of Complex Objects", *Proc. ECOOP/OOPSLA'90*, pp. 245-257, (1990).
- [MFH95] Mukherjea, S., Foley, J., Hudson, S., "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views", *Proc. ACM CHI'95 Conference*, pp. 331-337, (1995).
- [NWS86] Norman, K., Weldon, L., Shneiderman, B., "Cognitive layouts of windows and multiple screens for user interfaces", *Intl Journal of Man-Machine Studies*, 25, pp. 229-248, (August 1986).
- [NSP96] North, C., Shneiderman, B., Plaisant, C., "User Controlled Overviews of an Image Library: A Case Study of the Visible Human", *Proc. ACM Digital Libraries '96 Conference*, pp. 74-82, (1996).
- [PCH92] Plaisant, C., Carr, D., Hasegawa, H., "When an intermediate view matters: a 2D browser experiment", University of Maryland Computer Science Dept Technical Report #2980, (October 1992).
- [PCS95] Plaisant, C., Carr, D., Shneiderman, B., "Image browsers: taxonomy, guidelines, and informal specifications", *IEEE Software*, 12(2), pp. 21-32, (March 1995).
- [RLS96] Roth, S., Lucas, P., Senn, J., Gomberg, C., Burks, M., Stroffolino, P., Kolojehchick, J., Dumire, C., "Visage: A User Interface Environment for Exploring Information", *Proc. Information Visualization*, IEEE, pp. 3-12, (October 1996).
- [Shn96] Shneiderman, B., "The eyes have it: A task by data type taxonomy of information visualizations", *Proc. IEEE Visual Languages '96*, pp.336-343, (1996).
- [Shn98] Shneiderman, B., *Designing the User Interface: Strategies for Effective Human-Computer Interaction, Third Edition*, Addison-Wesley, (1998).
- [SSS86] Shneiderman, B., Shafer, P., Simon, R., Weldon, L., "Display strategies for program browsing: concepts and an experiment", *IEEE Software*, 3(3), pp. 7-15, (March 1986).
- [Twe97] Tweedie, L., "Characterizing Interactive Externalizations", *Proc. ACM CHI '97 Conference*, pp. 375-382, (1997).
- [Wil96] Wills, G., "Selection: 524,288 Ways to Say 'This is Interesting'", *Proc. Information Visualization '96*, IEEE, pp. 54-60, (1996).