

# Hierarchical Linked Views

Robert F. Erbacher  
Department of Computer Science, UMC 4205  
Utah State University  
Logan, UT 84321  
Robert.Erbacher@usu.edu

Deborah A. Frincke  
Pacific Northwest National Laboratories  
Richland, WA 99352  
Deborah.Frincke@pnl.gov

## ABSTRACT

*Coordinated views have proven critical to the development of effective visualization environments. This results from the fact that a single view or representation of the data cannot show all of the intricacies of a given data set. Additionally, users will often need to correlate more data parameters than can effectively be integrated into a single visual display. Typically, development of multiple-linked views results in an adhoc configuration of views and associated interactions. The hierarchical model we are proposing is geared towards more effective organization of such environments and the views they encompass. At the same time, this model can effectively integrate much of the prior work on interactive and visual frameworks.*

*Additionally, we expand the concept of views to incorporate perceptual views. This is related to the fact that visual displays can have information encoded at various levels of focus. Thus, a global view of the display provides overall trends of the data while focusing in on individual elements provides detailed specifics.*

*By integrating interaction and perception into a single model, we show how one impacts the other. Typically, interaction and perception are considered separately, however, when interaction is being considered at a fundamental level and allowed to direct/modify the visualization directly we must consider them simultaneously and how they impact one another.*

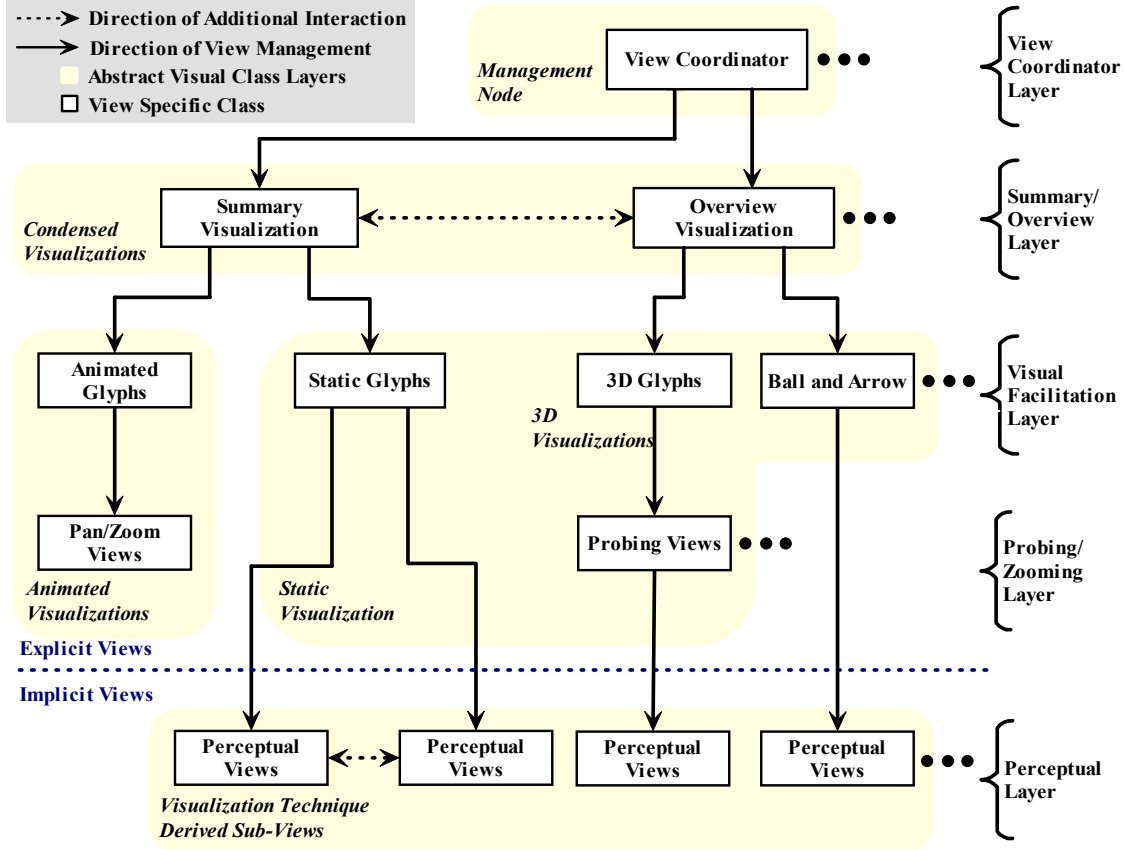
**KEYWORDS:** Visualization, perception, interaction, coordinated views, hierarchical model, intrusion detection

## 1. Introduction

Visualization has proven effective at the representation of disparate data sets to aid users in its analysis and comprehension. Users, however, will often need to correlate more data parameters or data elements than can effectively be integrated into a single visual display. Additionally, a complete analysis process generally requires the incorporation of multiple visualization techniques as well as varied views of the data, as no single view or representation of the data can show all of the intricacies or relationships inherent to a given data set. Thus, coordinated views have proven critical to the development of effective visualization environments.

Typically, development of multiple-linked views results in an adhoc configuration of views and associated interactions. In this paper we propose a hierarchical model geared towards more effectively organizing such environments and the views they encompass. In so doing we are providing guidelines such that implementers and researchers alike can better provide the infrastructural capabilities critical to making the most of the provided visualization techniques.

Our goal is to expand on prior art to create a model to assist in developing effective organizations of views, identifying consistent management schemes, and identifying where and how to implement linkages between views. We must ensure that the linkages are both effective and meaningful in order to be usable. In this fashion we propose a hierarchical model for the organization of views. In addition to the traditional types of views we also extend views to incorporate perceptual views. With perceptual views we are attempting to specify the multiple levels of foci incorporated into many visualization environments, as in [12]. We explicitly incorporate this concept of



**Figure 1: Hierarchical model overview.** Solid directed lines show directions of view management (control). These lines are also used to pass interaction, though in an undirected fashion. Dashed directed lines show additional interaction paths. Implicit views correspond to the different levels of focus applied by the user and the associated disparity in information gleaned from the image.

perception, including the impact that modifying the visualization techniques through interaction can have on the viewers' perception of visualizations.

In this way, a global (overall) focus within a visual display can provide one level of context, providing information on general trends and overall structure. Individual elements within the display, e.g., individual glyphs, can provide specific details with respect to an individual element. The user can thus change their focus from a wide view to a narrow view and ascertain differing qualities and values of the represented information. Thus, we have views within views differentiable by the user's perceptual focus (i.e., global versus local focus).

By interacting with the visualizations directly and changing their characteristics we are in essence changing the perception of the visualization. With the greater amount of direct manipulation being incorporated into visualization's we must consider

these elements of the visualization system as a whole rather than as separate entities as they are now.

## 2. Hierarchical Views

Traditionally, coordinated views have been treated as a series of sibling windows in which interactions within one window are carried through to other windows. This metaphor has been extended in Snap-Together Visualization [26], [27]. Snap-Together Visualizations allow the actual linkages to be more finely controlled. In this scenario, the views that are made available are dynamically controlled, as are the available linkages between these views and the interactions carried across them.

We propose an extension of this line of research to a full scale hierarchical model as shown in Figure 1. In this hierarchy, there can be any number of nodes (view instantiations) at each level. However, the relationship between the nodes will directly impact the

types of interaction passed between the nodes and the meaning or impact of that interaction on each view of the data.

The representative hierarchy shown in Figure 1 consists of five principle layers, including: the view coordination layer, the summary or overview layer, the visualization technique layer, the probing/zooming layer, and the perceptual view layer. Each layer presents an additional level of refinement in the view hierarchy. There may be any number of instantiated views in conjunction with each node of the hierarchy; instantiations are not shown.

The solid lines show direct connections in the implied top-down hierarchy. Dashed lines show additional possible routes of communication between views. The directed arrows show direction of control (i.e., creation and context management) and not flow of interaction, which may flow bi-directionally. The goal with the hierarchy is to present expected or comprehensible communication and control paths. While it is quite reasonable for all direct relatives to communicate, it would be unexpected for relatives to communicate without fully traversing the hierarchy. This would lead to interactions in which parent views become unsynchronized, resulting in inconsistent and confusing environments. This should not be confused with scenarios in which parent views cannot visibly represent an interaction; though representing the interaction at each stage aids in context and focus.

### **2.1. View Coordination Layer**

The view coordination layer is responsible for view creation, destruction, and management (e.g., showing and hiding views). This layer provides the global control of the view and linkage environment that cannot be effectively integrated within individual view layers.

### **2.2. Summary/Overview Layer**

The summary/overview layer encompasses visualizations designed to be representative of the database in its entirety; similar to the aggregate concept by Goldstein et al. [13]. Generally, this level of display will be far more coarse grained and thus require a significant loss of detail in comparison with its subviews, i.e. from the visual facilitation layer. The principal goal of this layer in the hierarchy is to provide context as to what is being represented within its subviews and to provide panning and region selection interaction facilities. Representation of selections made within a subview may not be

accurately represented in such displays due to the loss of detail and resolution.

While exact terminology may vary, we can conceive of multiple design goals for the visualizations at this level. With respect to overview visualizations we are conceiving of the detail visualization techniques being zoomed out such that an overview of the same data and visual characteristics are being represented. With summary visualizations we are considering completely separate techniques being used for the summary representation of the data; i.e., a technique specifically designed for the large-scale data en masse. In both cases the goal is to allow the representation of the entire data within a single window simultaneously, which may not be possible with the detailed views.

### **2.3. Visual Facilitation Layer**

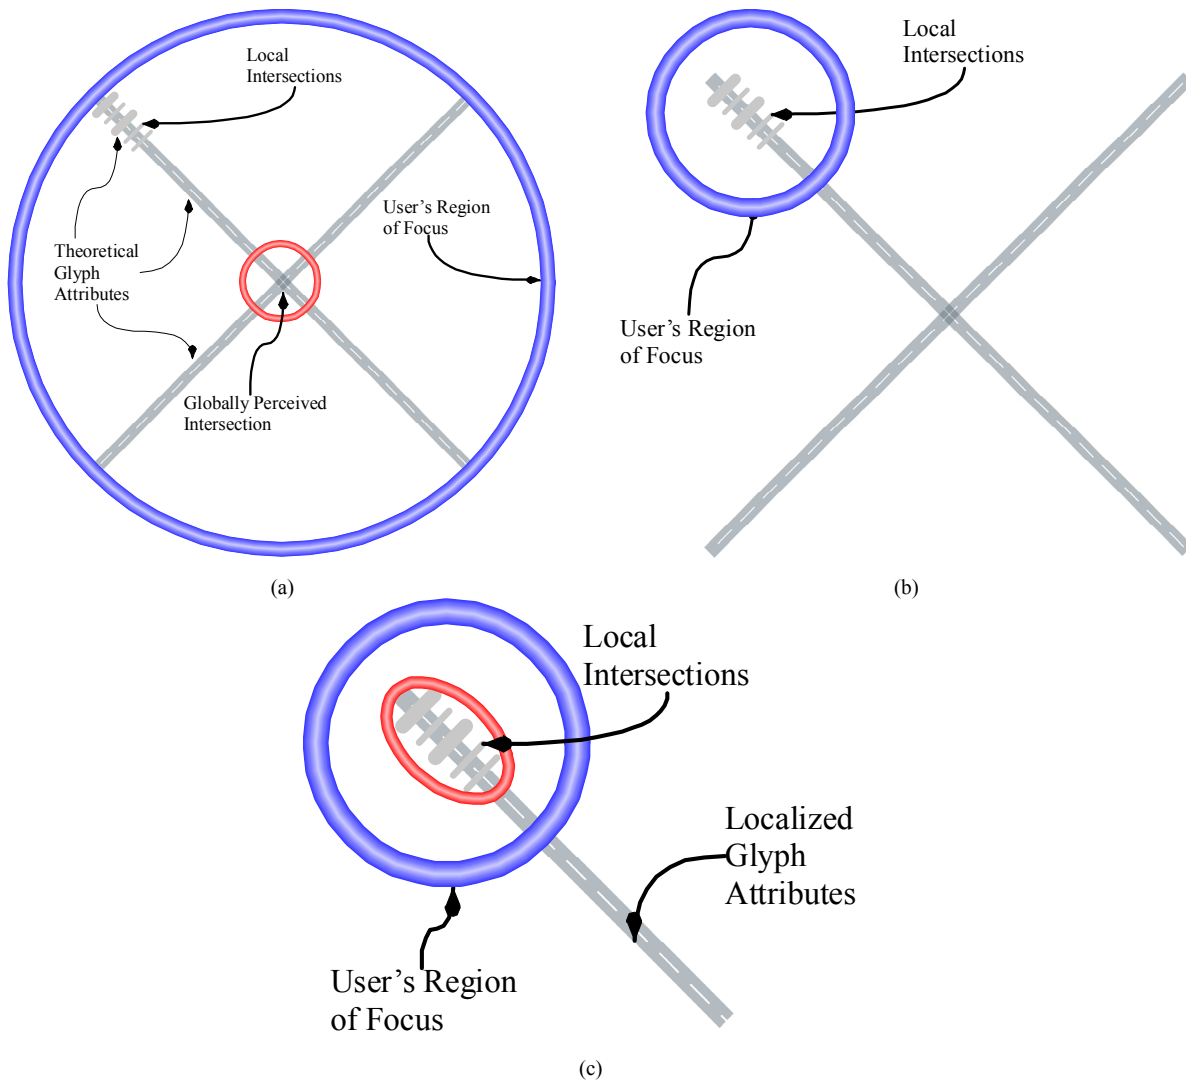
The visual facilitation layer incorporates the fundamental visualization techniques that make up the principal visualization architecture. This layer would incorporate all of the typical detailed views through which the user performs most of their analysis and interaction during the exploration and investigative processes. Typical techniques would include: animated 2D glyphs, static 2D glyphs, ball and arrow visualizations, graphs, space filling curves, trees, etc.

### **2.4. Probing/Zooming Layer**

The probing layer provides a detail or feedback layer. The information available through this layer is generally retrieved through direct interaction with the visualization techniques within the visual facilitation layer. The feedback can be provided through a separate interface panel, a separate window, or placed directly over or adjacent to the selected item (e.g., as an overlay).

### **2.5. Perceptually-Based Views Layer**

The final layer in our hierarchy is the perceptual view layer. This layer is based on the fundamentals of human perception and the way the human visual system aids users in the interpretation of images. Here we are relying on two processes. First is the process by which users will view the display in its entirety at a global scale before drilling down their focus to examine individual features or elements [21]. Second is the concept of perception of scale [34]. With perception of scale, larger elements will draw the user's attention before smaller details. This can be



**Figure 2: Examples of the impact of the user's focus, circled in blue. A global focus (a) results in perception of scale issues realizing the intersections circled in red. When the user drills in their focus (b) they resolve the more detailed intersections highlighted in (c).**

used in conjunction with the fact that crossed lines are perceived preattentively. Thus, a user interpreting Figure 2-a will first perceive the larger crossed lines and their associated intersection before drilling down their focus, Figure 2-b, to perceive the line intersections shown in Figure 2-c. This becomes a contrast between global or large-scale artifacts and local, detail specific artifacts.

These perceptual features can be used to essentially provide an additional level of detail, a new view of the data, within the same display, relying on the user's level of focus on the display. Thus the global view or focus will provide overall trends and general characteristics of the data while the drilled down focus will provide specifics on individual

elements. These individual elements could be represented as glyphs or icons.

Note that definition of all possible perceptual techniques is beyond the scope of this paper. Our intent is to provide a few examples of perception and how such views might be supported within our hierarchy. Multiple perceptual views could be included through the incorporation of different perceptual artifacts in which one group of characteristics are more quickly or readily perceived than a second group, e.g., preattentive versus non-preattentive attributes or first order versus second order statistically linked visual attributes [20].

Such perceptual views have disadvantages as well. Changes in the visual attributes of one

characteristic can readily change the perceived visual attributes of another characteristic. For example, it is well known that changing the background color changes the perception of the foreground color [20]. Similarly, highlighting or selecting one element may change the perception of another element. Thus, the perceptual views are implicitly linked; with interaction on one impacting the other.

## 2.6. Class Abstractions

The class abstractions in our proposed hierarchy are representative of the differences in interaction metaphors intrinsic to sets of visualizations and are indicative of necessary conversions. Ultimately, interaction classes are directly related to the type of visualization applied. This is shown in the examples provided within the regions highlighted in yellow; namely, animated visualizations, static visualizations, and three-dimensional visualizations. Clearly, the meaning of user interactions within each of these visualization paradigms will be vastly different and have vastly different meanings. This results from the difference in interpretation of dimensions and parameters between the abstractions as well as their application of spatial dimensions.

This does raise interesting issues as to how or even if some interactions should be converted. For example, a rotation within a three-dimensional view could either be ignored by linked two-dimensional views or converted into a rotation of the data parameter to axis mappings. The specifics of the relationship between those views will determine the most appropriate response.

The key to effective handling of interaction passing between different abstractions is in providing appropriate context. Therefore, when interactions are passed, it is important to pass sufficient information between the views with respect to the meaning and intent of the interaction such that each view receiving information about the interaction can correctly interpret the interaction according to its own paradigm. For instance, rather than merely passing a list of elements, pass the start and end indices as well as the operations performed and the context under which they were performed. Goldstein et al. have examined the issue of communication of interactions extensively [13]. Additionally, the reduction of interactions into their basic visualization interaction components [8] will reduce incorrect interpretations.

## 3. Interaction Linkages and Paradigms

In conjunction with the provided view hierarchy is the need to integrate an effective interaction paradigm. Given our interest in coordinated views we are clearly concerned with defining appropriate linkages to aid researchers in designing an effective interaction paradigm to be used in conjunction with the view hierarchy.

The guiding principle for the linkage metaphor in conjunction with the view hierarchy is to ensure the interactions follow the hierarchy. For example, an interaction in one view should be passed to that view's parent. The parent is then responsible for passing the information to its other children when appropriate. This ensures the parent can provide the necessary context for the interaction and maintain state information for each of its children, much of which will likely be visually represented.

Employing the described view hierarchy and linkages will result in scenarios of cascading interactions. In such scenarios, an interaction in one view is passed to its parent, which then passes the information to all of its children. This can then be applied recursively resulting in interactions being passed quite widely throughout the hierarchy. However, while it is important for the direct parent of a view to show activity within its children, it is less critical for this interaction to be passed up further in the hierarchy. Thus, implementations of such a hierarchy can allow more dynamic control over when and where to pass interactions beyond one level. The farther a node is in the instance hierarchy from the node interacted with, the less relevant that interaction likely is to the effective interpretation of that node's display. For instance, one of the benefits of having multiple levels in the hierarchy is to provide multiple levels of detail. The further down the hierarchy the greater the detail displayed. Conversely, going up the hierarchy leads to less detail and more abstractions. This reduced detail indicates that interactions more than one step away aren't necessarily viewable. Thus, while it should be configurable as to how far interactions are passed it becomes less useful the further up the hierarchy the interactions are passed.

The only direct linkages between sibling views are at the top of the hierarchy in which we may have multiple summary/overview displays and at the very bottom of the hierarchy in conjunction with the perceptual views. In the diagram exhibiting the model in Figure 1, solid lines identify paths of interaction passing. The interaction passing is assumed to be

bidirectional. The dashed lines identify additional paths of interaction passing.

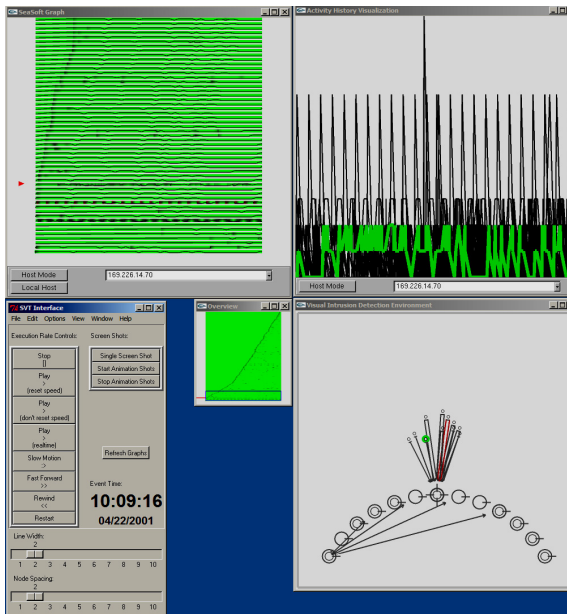
Linkages between the perceptual views are predicated on the fact that changing perceptual characteristics of the display at one level will impact the perceptual characteristics of the display at many levels. This results in implicit changes to the views rather than the explicit changes we have been considering previously. Given the extent to which direct manipulation is being applied for analysis—both to garner additional feedback and also to rapidly change visualizations, through the changing of mappings, display ranges, etc.—comprehending the impact of such interactions on the perceptual aspects of the display are critical. These impacts must then be taken into account when considering the visual analysis environment as a whole.

Finally, linkages must be used consistently. If an interaction linkage is identified between multiple views then this linkage should showcase all interactions. Interactions that don't make sense to incorporate in the linked view or are disabled should still provide feedback indicating what interactions have been filtered. We have not discussed the configurability of views and linkages in this paper; see North et al. [26], [27]. A refinement to this metaphor


is that all interactions should be passed up the hierarchy. The parent view, however, can be more selective as to what information to pass to its children without any loss of context or consistency. An example where this may be the case is if the parent view provides separate context representations for each of its children. Failing to provide such context will lead to confusion as users attempt to follow interactions between displays.

#### 4. Examples Using the Hierarchy

We have provided a hierarchical model for views and attempted to describe the fundamental applications of the model. Here, we use of an intrusion monitoring and detection visualization environment [10], [11] to illustrate the usefulness of the five-layer hierarchical model. This environment visually represents intrusion related data (see Figure 3), an application falling within our motivation of supporting large, disparate data sets. The intrusion monitoring and detection visualization environment currently supports analysis of host-based data, including: the system log file, the system last log, and system statistics. This combination of data provides information as to connections and disconnections, kernel and



**Figure 3: Example showing all windows of the environment simultaneously. The same host is selected in each display. This selection is automatically transferred to all windows when selected in one. Probing views are shown in two of the windows.**



**Figure 4: Summary/Overview of the intrusion detection visualization. Selected region highlights context in sub views. The blue rectangle highlights the region presented in the sub-view. The red hash exemplifies the selected element.**

application messages, alerts, system load, number of users, etc.

We begin by providing some guidelines to apply towards an implementation in conjunction with the actual examples. Next, we explain how each of the layers is used in the intrusion monitoring and detection environment. Finally, we consider analysis by the user within the environment.

#### 4.1. Summary/Overview Example


Figure 4 shows an example of an overview window from our intrusion detection environment. The overview window has a visible instantiation that indicates the estimated severity of events by host in the data set over time. Note that the severity of an event is indicated by its color, using a green  $\Rightarrow$  red scale. Red is the most severe with green being the least severe. This scale was chosen to match typical cues of good vs. bad. Since the majority of the activity is green it essentially becomes a background color and deviations can easily be seen; whether they are in the red or black range; essentially an element of perceptual discrimination. Individual hosts are on the vertical axis and time is on the horizontal axis. In Figure 4, the most severe events are dark green/dark red, with no bright reds.

This example highlights several interaction and visual feedback metaphors we have hinted at. First, a node selected in any view is marked with a red dash on the side of the display. The currently viewed hosts are highlighted within the blue box. This is a necessity given that we cannot provide a detailed view of all hosts within a single display at once, given the number of available hosts. As indicated, it is not uncommon for an overview view instantiation to lack the ability to display some of the details that can be shown in more detailed views of the visual facilitation layer.

This view now only displays the information, it also servers as an organizer. It is the link between the view coordinator and the lower levels, and it is how we are maintaining context and ensuring that interaction follows naturally through the hierarchy.

#### 4.2. Static Visualization Examples

Figure 5 shows a static visualization indicative of the visual facilitation layer. Essentially, this is a more detailed representation of Figure 4, the overview view instantiation. This view allows for the representation of far more host detail than can be provided with the overview display. However, the inclusion of the



**Figure 5: Static visualization showing the subset of data highlighted in the overview visualization. A selected host is shown through the presence of the red triangle. This is the same selected element exemplified in Figure 4.**

additional detail prevents the display of all hosts simultaneously.


This display provides several *perceptual* levels or views of information. At the global level the analyst can pick out the severity of events occurring on the host. At the more detailed view level, lines are drawn within the glyph for each host to be representative of the connecting host; inspired by Seesoft [9]. In essence, these lines appear as a histogram overlaid onto each host's glyphs (The horizontal green  $\Rightarrow$  red bars). This perceptual focus technique aids identification of whether the shown activity is all from the same remote host or from different remote hosts; this is important when identifying the meaning and severity of a sequence of events. Thus, we essentially have multiple views within the context of the single display with the visibility of each element dependent on the user's focus, whether global or local.

**Selection and Feedback Example.** The user may select hosts within this display, as indicated by the red triangle. This selection is then passed to the parent window, which passes it to the other visualization windows.

**Pan/Zoom View Example.** Figure 5 also provides an example of a pan/zoom capability. Since it is limited in the number of hosts that can be displayed, the user can pan through the remaining hosts. The

hosts currently viewed are represented in the parent window through the use of a blue region selector.

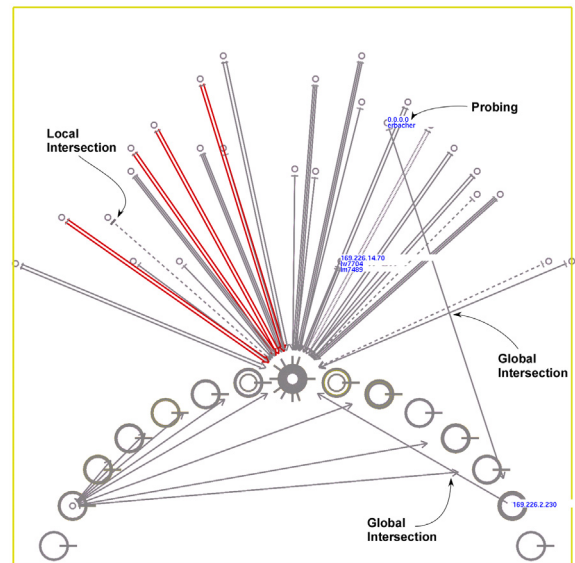
Figure 6 provides a second example of a static visualization example using the same underlying data set. This is a rudimentary display showing the severity level of activity as a histogram rather than a pixel plot. This view is a sibling of the one shown in Figure 5, still a member of the visual facilitation layer. This specific view is introduced as an example of selecting and highlighting nodes and how that supports aiding identification of activity in different displays. Alone, this histogram display is not very useful due to the amount of occlusion at the bottom of the display. Each individual element of the histogram, however, is more valuable than in the pixel plot since the histogram provides a clearer representation of activity over time. Used together with the histogram, the pixel plot provides an effective mechanism for selecting hosts while the histogram plot provides a more effective mechanism for representing individual elements. Thus, as coordinated views the whole of the environment is far more valuable than merely the sum of the parts, i.e., the individual views.



**Figure 6:** A second static visualization example is shown. The same host is selected/highlighted in this example as in Figure 5 (shown as a green selected line), showing the coordination between these two sibling displays and the parent. The importance of such coordination derives from the clarity of the selected host within the highly cluttered image.

### 4.3. Animated Visualization Example

Figure 7 shows an example of an animated visualization example, another example of the visual facilitation layer. The display shows a representation of network activity at one point in time. This example shows the most amount of detail but is limited to displaying the host's active at a single point in time. Clearly, this representation is completely distinct from the representation in the parent and sibling views. Currently, the environment is designed to show all hosts in the animated representation. An alternative representation would be to limit the represented hosts to those shown in the region selection. However, the region selection was designed with the static visualization technique in mind so as to mitigate that views limitation on the number of viewable hosts. This divergent philosophy is representative of the differing abstraction classes associated with the visualization techniques incorporated within each view. The differing abstraction classes make it clear that the interaction and feedback within one view will not necessarily make sense in the other views.



**Figure 7:** Animated intrusion-monitoring visualization. This display shows another view of the data in conjunction with probing as well as multiple levels of focus. At the global focus intersections and line angles provide the principal focal points, identifying anomalies. At the local level, intersections identify hash marks used to identify user volume; allowing retrieval of specifics.

While not shown, this visualization technique will maintain the representation of selected hosts by highlighting selected nodes in green. This provides consistency and allows for a host selected in one view to be quickly examined in all views. This also shows that while some interactions do not make sense across all view boundaries, some clearly do.

**Perceptual Views Example.** This visualization also provides an example of perceptual views. Specifically, each individual remote host glyph contains hash marks on its interconnection line indicating how many connections are being made. These hash marks provide a local perceptual feature, using line intersections to aid visual identification. However, these hash marks are small in scale, especially in comparison to the large-scale line intersections generated by anomalous activity. In Figure 7, a user connecting from a remote location to a local workstation, rather than a server, is considered anomalous and is highlighted by the numerous crossed lines, of large-scale, and the divergent line angle. These large-scale elements would be clearly picked up during immediate investigation of the image while the smaller scale features would be picked up during detailed analysis (i.e. scanning) to determine more precisely what is occurring and why.

**Probing and Feedback Example.** Figure 7 also provides examples of probing views, part of the probing/zooming layer. In this display, we have probed several elements within this display to retrieve additional information, which is provided as textual feedback directly adjacent to the picked elements. The information provided includes: the hostname, the IP address, and the usernames of those with active connections. This provides the detailed specifics needed to perform extensive analysis and follow up on anomalous activity detected within the environment.

#### 4.4. Analysis of Hierarchical Model Utility

The described environment benefits from the proposed hierarchy in additional ways. First, given the coordination of the children with the parent, gaining context and reorienting to the new display is far more efficient. Without the employed synchronization the environment might well be confusing and require substantial effort to maintain and re-establish context. For example, with the overview display, it would become very easily to disconnect and require effort to correctly associate the information shown in this display with the correct window to which it is associated.

With an increase in the number of instantiated views (e.g., multiple overview/summary displays) it becomes particularly important to show to which overview/summary display a window belongs. This is related to the way in which probing is provided in our environment. The integration of the detail views with their associated windows ensures that the detail view for one window cannot be confused with the detail view of another window.

The application of multiple levels of perceptual views allows for far more information to be integrated within a single display. This is exhibited in two of our displays, Figures 5 and 7. In these figures the user can identify artifacts at the global level and perceptually drill down to interpret more specifics. In Figure 7 the large crossed lines and deviating line angles stand out at a global level while number of users, hash marks, and connection information, line style, are more localized. Without such integrations the user would be required to apply for more interaction to garner the same level of information and comprehension.

Information is repeated in multiple displays. This is true of highlighting, which is shown in all displays, as well as visible elements which are highlighted in two of the displays. This repetition aids analysis and context. It is very easy to associate the elements displayed in the static representation of Figure 5 with the remainder of the available data elements due to the context provided by the summary display of Figure 4.

Additionally, context can easily be maintained from one display to another due to the repetition of selected node highlighting. A single display will not show all characteristics or relationships of a data set. Without the coordination of highlighting it would take an enormous amount of time to locate an element identified as being anomalous in other displays that would provide additional aid in the analysis of said element.

#### 5. Relation to Previous Work

The concept of views originated with databases to represent subsets of the available parameters to be viewed simultaneously. This greatly reduces the burden of examining a large database, especially given that many of today's databases contain hundreds if not thousands of parameters. In providing visual representations of such databases, visualization necessarily required a corresponding metaphor of views in order to similarly reduce the burden of parameter heavy databases. The visual metaphor of the view relegates itself to a display or window, i.e., a

different visual representation. As such, each visual representation may incorporate varying database parameters. Coordinated views, i.e., multiple linked views, extend the concept of views to transport interactions from one view to another to assist in correlation and exploratory practices. Ultimately, coordinated views have been found to compliment the focus and context [5] requirements of typical visualization environments.

Coordinated views have found currency in a wide variety of environments. For example: Mukherjea et al. discuss the application of views in conjunction with hypermedia networks to show relationships among different parameters within the database [24]. Gresh et al. discuss the use of views and the criticality of linkages in order to effectively examine protein simulations [14]. Gresh et al. also applied linked views to cardiac simulations in the WEAVE environment [15]. The authors point out the effectiveness of coordinated views to show how relationships are associated between views. This is done by selecting regions in one view and seeing the results in all associated views.

While numerous examples of coordinated views can be found in the literature, little work has been done to improve the science behind coordinated views. Boukhelifa et al. [3] provide the most relevant work in their application of coordination to multiple-view visualizations. In this work, the linkages share parameters and interactions; and not data values. This work can easily be applied on top of our hierarchical model.

Hill [22] enhanced typical view linkages by extending such linkages to incorporate constraints and deriving more efficient communication paradigms for linkages. Baldonado et al. provide an initial set of guidelines for when and how to use multiple views [1]. This work derived from their experiences at a workshop in '2000. North and Schneiderman [26], [27] developed snap-together visualizations to allow rapid view generation and linkages. They also performed the only user study to determine the effectiveness of coordinated views and their usability. Roberts [30], [31] formalized the data flow model characteristic of multiple views and has been aggressive in urging the effective use of coordinated views in visualization environments.

Taxonomies and models have been proposed for several related aspects of the visualization paradigm. Card et al. [4] provide an initial set of categorizations and associated differentiation parameters that signify the design characteristics incorporated into the

selected visualization techniques, the corresponding design space. Pfitzner et al. [28] provide a taxonomy for information visualization based on user skill level, user context, data type, task type, interactivity type, and visualization technique or type. Similarly, Chi et al. [6], [7] provide a visualization taxonomy based on the data state reference model. Chi's model is based on the current representation of the data and incorporated operators and interactions and their impact on the data state. Schneiderman [32] provides a related taxonomy based on seven different data types and seven interaction tasks, more specifically his work "offers a task by data type taxonomy with seven data types (1-, 2-, 3- dimensional data, temporal and multi-dimensional data, and tree and network data) and seven tasks (overview, zoom, filter, details-on-demand, relate, history, and extract)." Additionally, Tweedie [33] provides a classification scheme for interactions in conjunction with the visualization characteristics, namely the types of data supported and information characterized in its display.

Two of the more interesting taxonomies that are more closely related to our work is that by Bederson et al. with Pad++ [2], North et al. with their "A Taxonomy of Multiple Window Coordinations" [25], and Goldstein et al. with their interaction frameworks [13]. Pad++ is designed specifically around a zooming interaction metaphor, often within a single window, to more effectively allow the user to focus in on content of interest within dense information displays. This maps well to our concept of views within views; in essence a single window is being used to provide what requires multiple windows in most environments. North et al. describe a taxonomy for coordinated multiple windows that provides guidelines for the use and application of coordinated views. These taxonomies are important as they are the first steps and directly impact the development of models.

While the work by Goldstein et al. described an interaction framework for data exploration environments it falls short of identifying how multiple windows should be related to one another. They describe the concept of aggregate windows but do not adequately describe the relationship of such displays with the remainder of the environment. Here we develop a formal organization of displays and their relationships.

In terms of perception, this work builds off of the fundamental perceptually based techniques of Pickett and Grinstein [29] as well as the fundamental work on the analysis of perception [17], [18], [19], [23]. The goal of this work is to create a model that integrates

this prior work on perception into a unified model that incorporates perception, interaction, visualization, and coordinated views. While we discuss new ways to consider perception and the impact on perception by interaction we do not advance the fundamental research into perception.

## 6. Summary

We have developed a hierarchical model for the incorporation of multiple linked views within a visualization environment. In conjunction with this model, we have provided linkage paradigms for user interactions that will aid in providing an effective, consistent, and comprehensible interface strategy. Finally, we expanded the typical concept of views to include the idea of perceptual views in which multiple levels of information can be incorporated into a single display by relying on perceptual characteristics, each generating differing visual responses or response rates.

The described model builds on prior work, such as that by Goldstein et al. [13], and identifies a more formal organization and relationships of windows and interactions. These concepts can easily be integrated with existing interactive and display oriented frameworks to create a more complete, consistent, and usable environment. It is the careful integration of interaction, perception, coordinated views, and visualization that makes this model truly unique. Such an integrated model is critical in modern environments which are beginning to attempt to incorporate such diverse characteristics. Visual analytics exemplifies the need for such an integrated model.

We have solidified aspects of the hierarchy through an example that applied our intrusion-monitoring environment to the proposed hierarchy. This environment provides multiple disparate yet coordinated views to aid analysts in both monitoring a networked environment for intrusions and analysis of anomalies to aid understanding and resolution of said anomalies.

Additionally, we have explored the implications of differing types of visualizations and how they should be associated for the greatest benefit. This directly impacts interactions and how the associated interactions will also be interpreted.

## 7. Future Work

We have discussed the relationship of views to the original database concepts which instantiated them and have discussed a hierarchical model for views. However, we have not adequately linked the two

concepts. Our goal must be to integrate our current model and the associated interaction facilities into a database model such that we can interact with a database directly and access more of the facilities available through a full database system than is currently available in visualization environments.

Additionally, we must examine mechanisms for relating the view hierarchy to the user and available interaction linkages to the user. Most often this will be obvious. However, large hierarchies of views can contain sufficient complexity as to require more substantial and specific feedback.

## 8. Acknowledgements

The comments and suggestions made by the reviewers through the several iterations of this paper have been greatly appreciated and contributed to the strength of the paper you see here.

## References

- [1] M.Q.W. Baldonado, A. Woodruff, and A. Kuchinsky, "Guidelines for Using Multiple Views in Information Visualization," in *Proceedings of the ACM Conference on Advanced Visual Interfaces '00*, 2000, pp. 110-119.
- [2] J. Bedersen, and J. Hollan, "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics," in *Proceedings of UIST '94*, 1994, pp. 17-26.
- [3] N. Boukhelifa, J.C. Roberts, and P.J. Rodgers, "A Coordination Model for Exploratory Multi-View Visualization," in *Proceedings of the Coordinated & Multiple Views in Exploratory Visualization Conference*, 2003, pp. 76-85.
- [4] S.K. Card, and J. Mackinlay, "The structure of the information visualization design space," in *Proceedings of the Symposium on Information Visualization '97*, 1997, pp. 92-99.
- [5] S.K. Card, J.D. Mackinlay, and B. Schneiderman, *Readings in Information Visualization: Using Vision To Think*, S.K. Card, J.D. Mackinlay, and B. Schneiderman (Editors), Morgan Kaufmann Publishers, 1999, pp. 306-309.
- [6] E.H. Chi, "A Taxonomy of Visualization Techniques Using the Data State Reference Model," in *Proceedings of the Symposium on Information Visualization '00*, 2000, pp. 69-75.
- [7] E.H. Chi, and J.T. Riedl, "An Operator Interaction Framework for Visualization Systems," in *Proceedings of the Symposium on Information Visualization '98*, 1998, pp. 63-70.
- [8] M.C. Chuah, and S.F. Roth, "On the Semantics of Interactive Visualizations," in *Proceedings of the Symposium on Information Visualization '96*, 1996, pp. 29-36.

- [9] S.G. Eick, J.L. Steffen, and E.E. Summer, "Seesoft - A Tool for Visualizing Line Oriented Software Statistics," *Readings in Information Visualization: Using Vision To Think*, S.K. Card, J.D. Mackinlay, and B. Schneiderman (Editors), Morgan Kaufmann Publishers, 1999, pp. 419-430.
- [10] R.F. Erbacher, K.L. Walker, and D.A. Frincke, "Intrusion and Misuse Detection in Large-Scale Systems," *Computer Graphics and Applications*, Vol. 22, No. 1, 2002, pp. 38-48.
- [11] R.F. Erbacher, Z. Teng, and S. Pandit, "Multi-Node Monitoring and Intrusion Detection," in *Proceedings of the IASTED International Conference On Visualization, Imaging, and Image Processing*, 2002, pp. 720-725.
- [12] Q. Gao, "Visual knowledge representation based on perceptual organization," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, San Diego, CA, 1998, pp. 4524-4529.
- [13] J. Goldstein, S.F. Roth, J. Kolojejchick, and J. Mattis, "A Framework for Knowledge-based Interactive Data Exploration," *Journal of Visual Languages and Computing*, Vol. 5, 1994, pp. 339-363.
- [14] D.L. Gresh, F. Suits, and Y.Y. Sham, "Case Study: An Environment for Understanding Protein Simulations Using Game Graphics," in *Proceedings of the Visualization Conference '01*, 2001, pp. 445-448.
- [15] D.L. Gresh, B.E. Rogowitz, R.L. Winslow, D.F. Scollan, and C.K. Yung, "WEAVE: A System for Visually Linking 3-D and Statistical Visualizations," Applied to Cardiac Simulation and Measurement Data, in *Proceedings of the Visualization Conference '00*, 2000, pp. 489-492.
- [16] G.G. Grinstein and H. Levkowitz, "The Importance of Teaching Perception in Visualization Courses," in *Proceedings of the First Eurographics Workshop on Graphics and Visualization Education*, 1993.
- [17] C. G. Healey, "Fundamental Issues of Visual Perception for Effective Image Generation," In *SIGGRAPH 99 Course 6: Fundamental Issues of Visual Perception for Effective Image Generation*, 1999, pp. 1-42.
- [18] C. G. Healey, "Applications of Visual Perception in Computer Graphics," In *SIGGRAPH 98 Course 32: Applications of Visual Perception in Computer Graphics*, 1998, pp. 205-242.
- [19] C. G. Healey, "On the Use of Perceptual Cues and Data Mining for Effective Visualization of Scientific Datasets," In *Proceedings Graphics Interface '98*, 1998, pp. 177-184.
- [20] W.R. Hendee, *Cognitive Interpretation of Visual Signals*, W.H. Hendee and P.N.T. Wells (Editors), Springer-Verlag, 1997, pp. 149-175.
- [21] C.A. Kelsey, "Detection of Vision Information," *The Perception of Visual Information*, W.H. Hendee and P.N.T. Wells (Editors), Springer-Verlag, 1997, pp. 33-55.
- [22] R.D. Hill, "The Abstraction-Link-View Paradigm: Using Constraints to Connect User Interfaces to Applications," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems '92*, 1992, pp. 335-342.
- [23] G. Liu, C. G. Healey, and J. T. Enns, "Target Detection and Localization in Visual Search: A Dual Systems Perspective," *Perception & Psychophysics*, Vol. 65, No. 5, 2003, pp. 678-694.
- [24] S. Mukherjea, J.D. Foley, S. Hudson, "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views," *Proceedings of the SIGCHI conference on Human factors in computing systems '95*, 1995, pp. 331-337.
- [25] C. North, B. Shneiderman, "A Taxonomy of Multiple Window Coordinations," *University of Maryland, College Park, Dept of Computer Science Technical Report #CS-TR-3854*, 1997.
- [26] C. North, and B. Schneiderman, Snap-Together Visualization: Evaluating Coordination Usage and Construction, *University of Maryland, College Park, Dept of Computer Science Technical Report #CS-TR-4075*, 1999.
- [27] C. North, and B. Schneiderman, "Snap-Together Visualization: Coordinating Multiple Views to Explore Information," *University of Maryland Computer Science Dept. Technical Report #CS-TR-4020*, 1999.
- [28] D. Pfitzner, V. Hobbs, and D. Powers, "A Unified Taxonomic Framework for Information Visualization," In *Proceedings of the Australian Symposium on Information Visualisation '03*, 2003, pp. 57-66.
- [29] R. M. Pickett and G. Grinstein, "Iconographic Displays for Visualizing Multidimensional Data," in *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics '88*, 1988, pp 514-519.
- [30] J. Roberts, "Multiple-View and Multiform Visualization," *Proceedings of the SPIE Conference in Visual Data Exploration and Analysis VII*, 2000, pp. 176-185.
- [31] J. Roberts. "On Encouraging Multiple Views for Visualization," *Proceedings of the International Conference on Information Visualization*, 1998, pp. 8-14.
- [32] B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations," in *Proceedings of the IEEE Symposium on Visual Languages*, 1996, pp. 336-343.
- [33] L. Tweedie, "Characterizing Interactive Externalizations," *Proceedings of the SIGCHI conference on Human factors in computing systems '97*, 1997, pp. 375-382.
- [34] J.M. Wolfe, "Visual Attention," In *Seeing: Handbook of Perception and Cognition*, K.K. De Valois (Editor), Second Edition, Academic Press, 2000, pp. 335-386.