

Cocktailmaps: A Space-filling Visualization Method for Complex Communicating Systems

Christopher Ahlberg

Department of Computer Science, SSKKII, and IVEE Development AB
Chalmers University of Technology
S-412 96 Göteborg, Sweden
ahlberg@acm.org

ABSTRACT

Cocktailmaps is a visualization method for visualization of communicative behavior in complex communication systems such as human conversation, cocktail parties, parallel computers, and telecommunication networks. Cocktailmaps are space-filling in that they effectively utilize the available screen real estate to communicate properties such as what communicators dominate a communication over time, what topics are communicated, and how agents move between subcommunications. Cocktailmaps have been implemented utilizing the Information Visualization and Exploration Environment (IVEE) which provides users of cocktailmaps with interactive techniques such as zooming, panning, filtering, and details-on-demand.

Keywords

Information visualization, dynamic queries, cocktailmap, spoken communication

INTRODUCTION

Complex communicating systems

Complex communicating systems are plentiful and include examples such as:

- Human beings in discussions, especially complex situations where multiple discussions occur and people move between discussions – i.e. cocktail parties and lunchroom discussions.
- Parallel computers where multiple computations are processed and computations are dynamically allocated to processors.
- Telephone networks where multiple humans (or other levels of abstractions such as cities, countries or other regions) communicate.

Those interested in examining such systems, for example researchers in linguistics, programmers of parallel computers, and telephone network operators, face many difficult and interesting tasks. The tasks may involve agents differing from humans and computer processors, to telephone network nodes. However, many tasks are common to the different types of systems:

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

AVI '96, Gubbio Italy

© 1996 ACM 0-89791-834-7/96/05..\$3.50

- Who communicates with who?
- Who dominates the communication?
- What topics occur?
- What topics dominate the communication?
- Do some topics typically follow others?
- Do some specific sequences of agents exist – e.g. do some agents typically follow others?
- Does some specific part of the communication (a discussion or a computation) affect others, e.g. a does certain topic in the discussion at a lunch room table lead to others joining into that discussion?
- Do agents move between discussions?

Data collected for examination of phenomena such as these are typically very large and hard to overview and handle. They might be transcriptions of spoken language, databases of telecommunications data, or sample data from the communication in a multiple processor computer.

Visualization methods have proven successful in many areas of research and business. Typical tasks when exploring complex datasets, such as seeing trends, observing clusters and patterns, detecting anomalies and finding individual data elements, can all be greatly facilitated when utilizing visualization methods. Traditional methods obviously include approaches such as cartesian scatterplots, bar charts, and pie charts. With the arrival of interactive computer graphics systems many new approaches have been demonstrated lately, including the cone tree [11], the treemap [6], and the hyperbolic treebrowser for hierarchical data [7], the perspective wall for time dependent data [11], the table lens [10] and the starfield [1] for relational data, and the filter/flow metaphor for boolean queries [12].

Visualizations of communicating systems

A number of methods for visualization of communicating systems exist. A straightforward approach is a graph reflecting the physical layout of a network, showing the agents which are communicating as nodes and the fact that two agents are connected as a link between two nodes (Figure 1). This approach emphasizes the structure of the communication and users may observe for example clusters in the communication quite easily. However it may also be quite difficult to map quantitative variables onto the graph, such as how much communication that occurs and for how long it occurs. Further, a graph representation emphasizes a static view of the communication, while communication is inherently dynamic in its nature.




Figure 1. Sample graph visualizing a simple communicating system.

Another approach frequently used in the telecommunication industry is so called Message Sequence Charts (MSCs) [3]. A sample MSC can be found in figure 2. Each of the lines A, B, and C correspond to a communicating agent, for example a process in a computer. An arrow between two vertical lines X and Y correspond to X communicating with Y. The vertical direction also corresponds to time, and accordingly the lines representing communication may be angled to indicate that the communication is not carried out immediately.

The MSC approach emphasizes the dynamic nature of the communication, i.e. it deals with one of the problems with the straightforward graph. However it is still difficult to map quantitative variables on the MSC.




Figure 2. Sample Message Sequence Chart (MSC).

Yet another approach to visualizing communicating systems is cluster visualizations. By utilizing statistical clustering techniques, a two-dimensional spatial layout of communication patterns can be computed [4]. The cluster configuration in figure 3 shows the "centrality" of 11 speakers based on a who-follows-whom matrix of a discourse. The background shading reflects the total amount of utterances made by the corresponding speaker in the foreground. The most frequent speaker sequences are indicated with arrows proportional to the frequency.

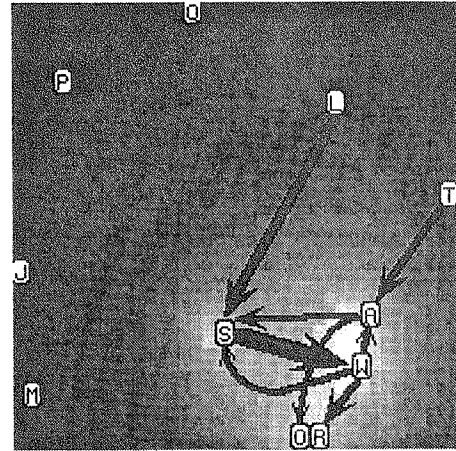


Figure 3. Cluster visualization of a group of people communicating.

COCKTAILMAPS

Cocktailmaps is an attempt to create a visualization method for complex communicating systems with the following properties and qualities:

- It should emphasize the dynamic nature of communication and clearly indicate the flow of time through the interaction between the participating agents.
- Quantitative variables such as the length of the communication and how much of the total a certain part of the communication occupies should be possible to map graphically to the visualization.
- Multiple agents communicating simultaneously, possibly with several subgroups, should be possible to map onto the same visualization, so that patterns where agents drifts between discussions can be examined.

Cocktailmaps originates from joint work between researchers in linguistics and computer science on how the interaction between multiple humans in a discussion with possible many subdiscussions, e.g. a cocktail party, can be graphically and visually represented.

A first example of cocktailmaps is shown in figure 4. This simple example visualizes the communication between four people, A, B, C, and D. Each of the three main columns correspond to the main sections of their conversation, i.e. time flows from left to right in the figure. The width of the columns is proportional to the length of the section compared to the total length of the conversation. The partitioning of a textual material into suitable sections could be achieved using various methods, including partitioning by hand, as for example done by researchers in linguistics who are studying spoken language, or by some automatic partitioning algorithm, for example as demonstrated by Hearst's textiling algorithm [5].

The rectangles in the columns correspond to the persons involved in the conversation, one rectangle per person involved in a section. The height of the rectangle is in this case proportional to how much the person verbally

dominates the discussion, as defined by number of words spoken by the person divided by the total number of words said in the section. Such a measure of dominance might be somewhat doubtful, but is one useful raw indicator of dominance in a discourse. The rectangles are in this case color coded indicating which persons are talking.

This space-filling approach to visualizing the communication between a number of agents, utilizes effectively the available screen space to show a number of variables, including who are speaking, for how long, how much, sequences of speakers, etc. This initial example demonstrates how cocktailmaps can be utilized for human communication – but many other variations are possible such as processors in a parallel computer or nodes in a telecommunication network. Below more complex examples of cocktailmaps will be discussed.

Origins of cocktailmaps

The origins of cocktailmaps can be found in several places, including treemaps [6], divided bar charts, and maybe primarily the Histomap [9]. The Histomap displays in a single visualization the history of world civilization (Figure 5). Cultures and nations are represented as colored areas along a time line. Several cultures are presented in parallel in a space-filling manner. The height of the visualization at any point of time is divided proportionally to the "power" of the respective cultures. Obviously this "power" differs over time, and the overall effect is a visualization where users

may observe how cultures grow and shrink over time, in the context of other cultures and nations declining and rising.

While the measure of "power" might be somewhat doubtful, the Histomap provides users with an intriguing and compelling overview of the world history.




Figure 5. Mock-up of a Histomap [9]. Each colored area corresponds to a nation and the area at any given time corresponds to the relative power of the nation/culture.

Treemaps are also a space-filling technique, similarly to the Histomap. Treemaps utilizes a slice and dice technique to display large hierarchies in a single screen. Each subdivision of the screen corresponds to a node in a hierarchy, and is given size dependent on a node attribute. Although the data




Figure 4. Simple cocktailmap example.

for which treemaps are utilized are different to those for cocktailmaps, i.e. hierarchies of various kinds, the actual screen layout is quite similar for treemaps and cocktailmaps.




Figure 6. A very small treemap holding only six nodes. Corresponding traditional tree layout displayed below.

Divided bar charts is a common approach to displaying statistical data, especially when the aim is to convey accurate impressions of parts of a whole. In a multiple divided bar chart, the length of each segment within a bar represents the proportion of one component of the whole, such as in figure 7. Cocktailmaps share some properties with multiple divided bar charts, in that cocktailmaps hold a number of columns which are divided into subrectangles which in turn are sized proportionally to some variable. In cocktailmaps the columns are placed bordering each other to emphasize the fact that the columns together form a flow of time. Cocktailmaps as they are discussed below also have interactive properties such as zooming abilities, details-on-demand, and filtering with dynamic query filters [1].




Figure 7. Divided bar chart example, showing the amount of the total export for three products for each of the countries Sweden, Denmark, and Norway.

COCKTAILMAPS AND IVEE

Cocktailmaps have primarily been applied to the area of visualizing spoken interaction between humans. The screen dumps of cocktailmaps in this paper are from an implementation created with the Information Visualization & Exploration Environment, IVEE [2]. IVEE¹ is an environment for automatic creation of *dynamic queries* environments where users may filter, highlight, zoom, pan and select details-on-demand in visualizations of large datasets.

Dynamic queries is a technique based on direct manipulation of database queries where users manipulate query devices such as rangesliders and toggles to interactively filter the contents of information visualizations such as geographic maps and starfields [1]. A typical dynamic queries example is the filmfinder where users can manipulate alphasliders to select actors, actresses and directors, a rangeslider to select film length, toggles to select PG-ratings and film categories, etc., and query results are presented in a starfield immediately upon user actions. By clicking on individual films in the starfield display, details are presented in a pop-up window, including both textual and graphical information (Figure 8).

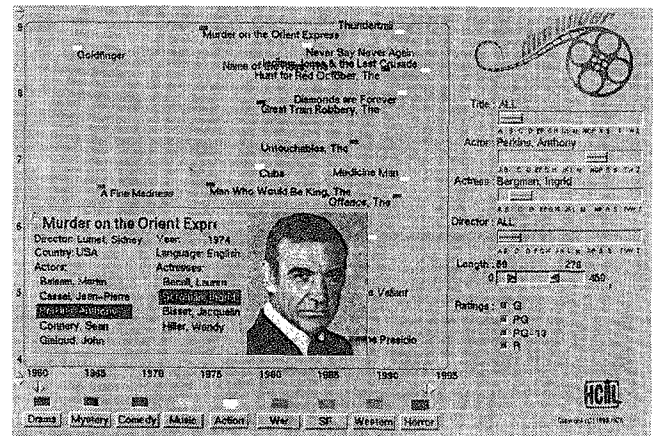


Figure 8. The filmfinder prototype [1].

IVEE allows users to automatically create dynamic queries environments such as the filmfinder [which was originally a specialized prototype just for a specific film database). IVEE examines a user specified datafile and automatically creates query devices and visualizations. While the basic visualization is the starfield, i.e. an interactive scatterplot which can be zoomed, panned, etc., users can create more advanced and complex visualization by supplying IVEE with files holding sets of polygons which are attached to database objects.

The simple cocktailmap in figure 4 was created utilizing this feature of IVEE. The database consists of a number of

1. IVEE is available on ftp at <ftp://ftp.cs.chalmers.se/pub/IVEE> or on the world wide web at <http://www.cs.chalmers.se/se/SSKKII/ivee.html>.

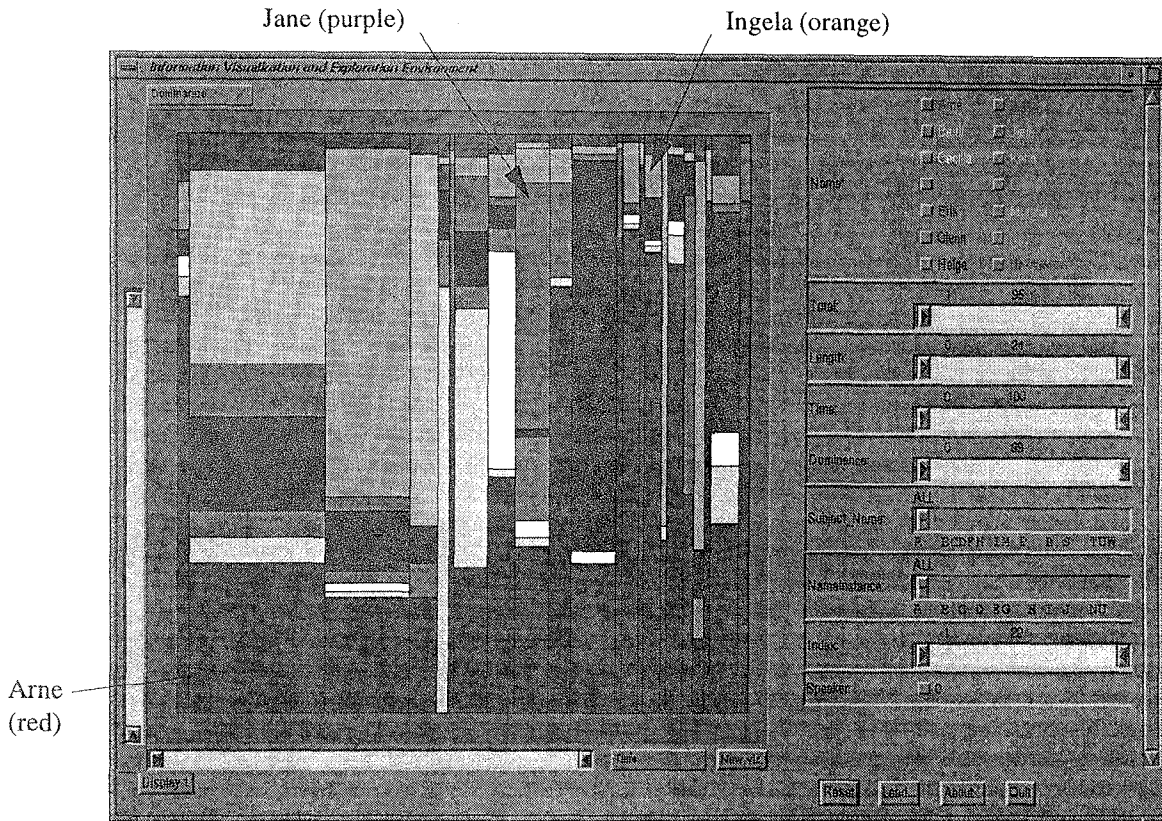


Figure 9. Cocktailmap where a meeting with fourteen participants is visualized.




Figure 10. Cocktail map which has been zoomed in vertical direction, three speakers have been highlighted, and details on a subsection has been selected.

tuples, each representing a person in a section of the conversation with attached variables such as length of the section, how much the person talks, the topic of the conversation etc. Along with this database, IVEE is supplied with a file of graphical objects, which are all rectangles, proportional in size to the appropriate database attributes.

A more complex example of cocktailmaps is presented in figure 9. This cocktailmap visualizes a meeting in a local town district board with fourteen participants and an agenda of 22 items. In the cocktailmap in figure 9 each of the participants has been assigned a color, so that each of them can be followed over time throughout the conversation. The group of toggles and labels on the top right each correspond to a speaker and in this case also indicates mappings between speakers and colors. The sections, i.e. the columns, correspond to each of the main items in the agenda of the meeting. This means that in this case the sections do not correspond to natural changes in topic in a conversation, instead they correspond to constraints set by the context.

As the cocktailmap is created inside IVEE, interactive operations such as zooming, filtering, and details-on-demand are provided for free. In figure 10 the cocktail map has been zoomed in the horizontal direction, i.e. time, to show the later part of the meeting in greater detail. The zooming is performed by dragging the zoombar at the bottom of the visualization. To the right in figure 9 and figure 10 a set of query devices is located, one for each database attribute.

By manipulating these query devices, users may filter the visualization to avoid the somewhat crowded display in figure 9. In figure 10, three persons have been selected and the subsections corresponding to them are left in the context of the structure of the whole visualizations. Finally details-on-demand has been selected on the speaker Jane by clicking on one of the subsection rectangles. The information is presented in the details-on-demand window on the bottom left in figure 10.

Other color schemes are obviously possible, for example coloring the subsections by which topic they cover. For the presented example this might not be so interesting as all the persons in the same section are assigned to the same topic in the database, i.e. the title of the item on the agenda. In other cases, for example when two subdiscussions are co-occurring this could be very useful.

COCKTAILMAPS IN CONTEXT

Below it will be demonstrated how cocktailmaps can be utilized for exploration of datasets collected from complex communicating systems. The main example discussed above, communication in a town district board meeting will be explored in depth for a number of different tasks. The data originates from a recording of a full meeting with the board, which has been transcribed into a computer readable format. A filter program was then created to transform the transcription into an IVEE readable database. The specific dataset might in some sense not be perfect for the purposes of this paper, for example the transcription does not indicate subdiscussions between participants in sections (columns) – the formal aspect of the conversation indicates that every

speaker in a section speaks to everybody else. The coloring of the cocktailmap in figure 9 is changed in figure 11 from coloring by person to coloring by topic.




Figure 11. Cocktailmap of the same data shown in figure 9 and figure 10, but colored by topic.

The meeting was four hours in length and the corresponding transcription is 5200 lines of text long, which obviously is very hard to overview and penetrate. The hope is that cocktailmaps can provide the overview needed to start exploring the complex data and the possibilities to examine quantitative details of sections and persons to gain deeper understanding.

The examples from human communication are complemented by a few examples from the domain of parallel computers.

Who communicates with who?

A very basic task for a dataset such as this is simply: who communicates with who, and how does it differ over time. In figure 9 it can be observed that the person corresponding to the red color (Arne), is continuously part of the conversation, which is confirmed when the actual transcription is studied – Arne is the chairman of the meeting. Some persons continuously take part in the discussion, for example Jane and Ingela (corresponding to the purple and orange colored subsections), while others just join in at certain times. These kind of patterns may in this context be due to personality type of the speakers (some people join in continuously in discussions while others are more quiet) or due to persons specific expertise in a certain topics.

Who dominates the communication?

The height of each subsection, i.e. rectangle in a column, is proportional to the number of words spoken by a particular person in that section. This measure can be interpreted as a dominance in the discussion, although it might not always be an altogether correct measure. We can easily identify for example how the chairman (red rectangles) verbally dominates the meeting throughout nearly every item of the agenda – sometimes the chairman is nearly the only speaker.

Some dominate temporarily, as discussed above, such as the large gray and blue rectangles, while others dominate in the

sense of being continuously part of the discussion, while never taking over completely. The width of the rectangles is important, i.e. for how long the measurement of dominance is valid. While the chairman Arne dominates some items of the agenda completely, these are not very long in terms of time, and in fact, are of the type "the chairman informs".

What topics occur and what topics dominate the communication?

As discussed, the specific example used here is somewhat rigid in that the topic of the transcribed discussion only change as the current item of the agenda changes. Below a variation of this example is displayed, where several subdiscussions and topics take place concurrently (Figure 12). The subsections are colored by the topic the person corresponding to that rectangle cover. While the basis for this example still is the transcription from a town district board meeting, the topics of discussions have been altered to demonstrate how cocktailmaps can be used to observe changes in topics throughout a discussion.




Figure 12. Cocktailmap displaying a communication pattern between a number of persons, with several subdiscussions. Subsections colored by topic.

In figure 12 five topics are covered in the visualized communication. Initially three topics are discussed with the one mapped to red dominating the communication space. As time moves on, more and more participants shift to the topic mapped to yellow, and for a short while all participants take part in a discussion on this topic. Finally a new topic is introduced which attracts several persons in the discussion, i.e. the blue rectangles, and an earlier topic, the green one, is introduced again.

Do some topics and persons typically follow others?

From a cocktailmap such as the one presented in figure 12 conclusions could be drawn about sequences of topics, such as how a the introduction of a certain topic will shortly lead to another being introduced. Typically phenomena such as this would be examined in several cocktailmaps visualizing different discussions to be able to draw conclusions about how certain topics affect a discussion. A simpler variation of this theme is how the chairman of a meeting typically initializes the discussion on each item in the agenda.

However, patterns of how certain topics initiate other topics might be even more interesting in other domains. The following example is from the domain of parallel computers. One processor might broadcast a message to a few other processors which instructs them to ask their neighbors to take part in the computation they are currently involved in. The processors transmit this message to their neighbors and they take appropriate action. How this kind of communication pattern might be visualized in a cocktailmap is shown in figure 13.

The communication over the communication bus is visualized (supposing the computer is a bus based parallel computer) with coloring indicating the type of message passed onto the communication bus. Each column correspond to various processors communicating over the bus, sampled at some appropriate time scale. The area of each subrectangle correspond to the bandwidth utilized by each communicating processor on the bus. The second column correspond to a single processor broadcasting a message to three processors. The message instruct these three processors to broadcast to their neighbors to take part in their current computation, displayed in column three. Finally in column four, these neighbors have started working on the same computations and communicate with their neighbors.




Figure 13. Mock-up of cocktail map showing communication in a parallel computer.

Do agents move between discussions?

The patterns in figure 12 indicated that persons moved between subdiscussions. This is a typical behavior for example at a cocktail party where some people move around rapidly only to shortly join discussions, while others stick with a fixed group. In a cocktailmap such as the one is figure 12 this is revealed by how a dominant person, shown by a large rectangle in a column, pulls the other persons into a certain discussion.

Exploring aspects of domination

The measure of domination utilized above is the number of words uttered by a particular speaker in a certain section. Some interesting observations can be made, in both the human communication and the parallel computer cases by complementing that measure with the number of sentences

uttered by the speaker. In the case of human communication this might indicate the difference between for example talking for a long while a few times and talking shortly a lot of times. In figure 14 the height of the subsections is still mapped to dominance (as defined by number of words uttered) and the share of sentences uttered in a section is mapped to the color of the rectangles, ranging from dark blue to white (from low to high). This allows for detection of patterns such as:

- The persons dominating in terms of number of words in section three dominate more in sense of the number of sentences, than the persons dominating in terms of number of words in section two.
- In section fourteen only two speakers participate, Jane and Arne. While they utter equally many sentences, Arne totally dominates in terms of the number of words.

In terms of the parallel computer example above, this approach to coloring could reveal interesting patterns. A typical task when examining parallel computer systems is trying to understand where the bottlenecks are, in communication or computation. This task can be approached by letting the area of the subrectangles still indicate how much of the available bandwidth on the communication bus is occupied by each processor, and letting the coloring indicate how much of the possible computational power a particular processor utilizes. Using that color schemes, processors which are communicating a lot, but performing little computation and vice versa can be revealed.

CONCLUSIONS & FUTURE RESEARCH

Cocktailmaps is a novel visualization method for visualizing communicative behavior. Cocktailmaps have been applied to visualization of spoken language in cooperation with researchers in linguistics. Other potential areas of use for cocktailmaps are visualization of network traffic in telecommunication networks and visualization of communication in a parallel computer. Many areas for future research are available, such as:

- Scaling of the technique. How do cocktailmaps scale? The example in figure 9 holds 22 sections with approximately seven agents communicating in each section. In many communicating systems both the number of agents and the time span is much larger. One approach to dealing with such scaling problems is a hierarchical decomposition of the visualization where users initially are shown an overview and by clicking on for example a section, that section is split into subsections with a smaller granularity in time. Another approach is clustering in various ways, where groups of communicators are represented as a single graphical object which in some appropriate way catches the properties of the group.
- The presented examples of cocktailmaps display data by position, width, height, and color of the rectangles. Other possibilities include labeling to indicate who are speaking (this was not included in the above examples due to technical limitations of the tool IVEE), arrows which indicate who speaks to whom in a section, other graphical symbols in subsection rectangles which can indicate




Figure 14. Cocktailmap where the height of the subsections is still mapped to dominance (as defined by the number of words uttered) and the number of sentences uttered is mapped to the color of the rectangles, ranging from dark blue to white (from low to high).

various properties, etc.

- Less strict sections. The slicing of the space into section does not need to be as strict as demonstrated in this paper. In a less formal conversation for example (as compared to the meeting discussed above) people will join in and leave a discussion quite arbitrarily, and this can nicely be reflected in the cocktailmap by just sampling speakers at a finer granularity and allowing subsection rectangles to overlap.

ACKNOWLEDGMENTS

The idea of cocktailmaps grew out a series of discussions on visualization of spoken language with researchers in computer science and linguistics, including Staffan Truvé, Erik Wistrand, Jens Allwood, Johan Hagman, and Lars-Erik Holmqvist.

REFERENCES

1. Ahlberg, C., Shneiderman, B. (1994), Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. *Proceedings ACM CHI'94: Human Factors in Comp. Systems*, pages 313-317. Also in Baecker, R., Grudin J., Buxton, W., and Greenberg, S., *Readings in Human-Computer Interaction: Toward the Year 2000 (2nd Edition)*, Morgan Kaufmann Publishers, San Francisco, CA, 1995.
2. Ahlberg, C., Wistrand, E., (1995), IVEE: An Information Visualization & Exploration Environment, *Proceedings of IEEE Symposium on Information Visualization Info-Vis'95*, Atlanta, 1995.
3. *CCITT Z.100 Standard* (1989), International Telecommunication Union, Geneva 1989, ISBN 92-61-03751-8.
4. Hagman, J. (1995), Chains of Brains - Statistical Analysis of Speaker Order, SSKKII Technical Report (forthcoming), Göteborg University, 1995.
5. Hearst, M. (1995), TileBars: Visualization of Term Distribution Information in Full Text Information Access, *Proceedings ACM CHI'90 Human Factors in Comp. Systems*, pages 59-66.
6. Johnson, B., Shneiderman, B. (1991), Tree-maps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures, *Proceedings IEEE Visualization'91*, pages 284-291.
7. Lamping, J., Rao, R., Pirolli, P., A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies, *Proceedings of ACM CHI'95: Human Factors in Computing Systems*, pages 401-408.
8. Mackinlay, J., Robertson, G., Card, S. (1991), The Perspective Wall: Detail and Context Smoothly Integrated, *Proceedings of ACM CHI'91: Human Factors in Computing Systems*, pages 173-179.
9. Rand McNally, *Histomap of World History* by John B. Sparks, ISBN 0-528-83426-6.
10. Rao, R., Card, S. (1994), The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information, *Proceedings of ACM CHI'94: Conference on Human Factors in Computing Systems*, pages 318-322
11. Robertson, G., Mackinlay, J., Card, S. (1991) Cone Trees: Animated 3D Visualizations of Hierarchical Information, *Proceedings of ACM CHI'91: Human Factors in Computing Systems*, pages 189-194.
12. Young, D., Shneiderman, B. (1993) A graphical filter/flow model for boolean queries: An implementation and experiment, *Journal of the American Society for Information Science* 44(4), pages 327-339.