

# TrustNeighborhoods: Visualizing Trust in Distributed File Sharing Systems

N. Elmqvist and P. Tsigas

Chalmers University of Technology, Sweden

---

## Abstract


*We present TrustNeighborhoods, a security trust visualization for situational awareness on the Internet aimed at novice and intermediate users of a distributed file sharing system. The TrustNeighborhoods technique uses the metaphor of a multi-layered city or fortress to intuitively represent trust as a simple geographic relation. The visualization uses a radial space-filling layout; there is a 2D mode for editing and configuration, as well as a 3D mode for exploration and overview. In addition, the 3D mode supports a simple animated “fly-to” command that is intended to show the user the context and trust of a particular document by zooming in on the document and its immediate neighborhood in the 3D city. The visualization is intended for integration into an existing desktop environment, connecting to the distributed file sharing mechanisms of the environment and non-obtrusively displaying a 3D orientation animation in the background for any file being accessed over the network. A formal user study shows that the technique supports significantly higher trust assignment accuracy than manual trust assignment at the cost of only a minor time investment.*

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Systems]: User Interfaces I.3 [Computer Methodologies]: Computer Graphics H.5.1 [Information Systems]: Multimedia Information Systems

---

## 1. Introduction

Distributed file sharing systems are now commonplace in today’s Internet-connected society and are a great way for users across the world to share and exchange information between each other. However, it is clear that security is a vital aspect for this kind of file sharing to succeed. As we continue to blur the border between local files and remote files on the network, it is becoming increasingly important to categorize files according to their *trust*, a function of the average trust of their owners as well as the owner’s own classification of the file. To compound this problem, the majority of the intended users for the new generation of file sharing systems do not necessarily possess a high level of technical knowledge, and can be seen as novice or intermediate computer users. Most of these users lack a clear mental model of computer security. It is clear that we must find ways to make the concepts of trust



**Figure 1:** *TrustNeighborhood visualization for a network of 2000 hosts.*

and security explicit even to such a relatively inexperienced audience.

In this paper we present TrustNeighborhoods (see Figure 1), our attempt at addressing this issue through the use of information visualization. TrustNeighbor-

hoods is a tool for graphical representation of document trust relationships in large-scale distributed file sharing systems. It is intended to convey a tangible mental model to the user and improve his or her security situational awareness on the Internet. Based on the model of human trust presented in Ben Shneiderman's treatment of "circles of relationship" [Shn02], TrustNeighborhoods uses the metaphor of the network being represented by a layered city (or fortress) with individual users or documents visualized as buildings and organized into geographical regions representing different trust intervals. Each region of the city is color-coded and visualized as a concentric ring; entities are automatically laid out in each ring with their size, color, and orientation conveying information about the trust properties of each entity. The 2D mode of the visualization allows for interaction with the trust model, including dragging-and-dropping documents and users to change their trust levels, as well as radial grid distortions for the purpose of studying specific parts of the dataset. The 3D mode, on the other hand, affords overview and navigation of the distributed file system, drawing each entity as a three-dimensional building and supporting the user metaphor of trust in the context of a city. The 3D mode also has a special "fly-to" command that smoothly animates the view from an overview to the position of the requested entity, showing its position in the trust relationship as well as its immediate neighborhood.

We have developed a prototype implementation of the TrustNeighborhoods visualization technique using a simulated distributed file sharing system; the input data is static and derived from an XML file, but the interface with the data source is exchangeable with a real file sharing implementation. Using this prototype, we have also conducted a formal user evaluation of the technique involving twenty subjects recruited from our university. Our analysis shows a significant improvement in trust assignment accuracy for our technique in comparison to manual trust assignment. While this comes at the cost of some additional time spent on classifying an entity, this initial investment will significantly speed up subsequent accesses by adding the entity to the user's known universe.

## 2. Related Work

The human factors aspect of computer security is often overlooked; even if a program is secure, security may still fail if it is used improperly. Whitten and Tyger [WT99] notes that more than 90% of all computer security failures happen due to configuration errors, facts that indicate that security is inescapably an user interface design problem. They go on to analyze

the PGP privacy software to point out examples of inadequate design that may provoke users to perform fatal mistakes, such as sending unencrypted messages or divulging private information. Yee [Yee02] argues that usability and security need not conflict, and presents a number of general design principles for designing secure and usable software. For file sharing systems, studies show that users of such systems often have difficulties understanding which of their own files are shared and which are not, and many unwittingly share personal or private files on the public network [GK03]. Our work is an application of these ideas, attempting to provide a usable visualization of security to facilitate secure file sharing.

The security concept we employ in this treatment is the concept of *trust*. Going into a detailed description of related work in this field is outside the scope of this paper; please refer to Marsh [Mar94] for references to other work in the field. Rather, in this paper, we concern ourselves with the visualization of trust using techniques from the information visualization area. Not many prior examples of trust visualizations exist; the most straightforward approach is to show trust as a node-link diagram, one such example being the visualization of decentralized "webs of trust" for PGP keys using the sig2dot [sig05] and similar programs.

The TrustNeighborhoods visualization draws influences from a number of sources; it is a radial space-filling (RSF) technique akin to [AH98, Chu98, SZ00], but it is not used for hierarchy visualization (like in the cited cases) and thus does not possess the parent-child property for circle arcs of classic RSF techniques. The interaction technique for radial distortion of rings is a focus+context [Fur86] technique similar to that of the InterRing [YWR02] system, also an RSF technique. Many examples of hierarchy visualizations that could be used for file systems, both local and distributed, exist; for example, the classic cone trees [RMC91], hyperbolic layouts [LR96, LRP95, MB95], and botanical tree visualizations [KvdWvW01], to name just a few (see Stasko et al. [SCGM00] for a survey). However, objects in distributed file sharing networks are typically organized in flat and shallow hierarchies, and thus the focus of the TrustNeighborhoods technique lies not on scalable hierarchy visualization, but rather on the user's cognitive model of security and trust. The Tudumi [TK02] system is another security visualization tool based on concentric circles, but targets the representation of computer logs.

## 3. Trust Visualization

In our adaptation of Shneiderman's circles of relationship, we map his ideas to an even more tangible metaphor: a multi-level city (or fortress) used as a

model for trust relationships. The intuition is the geographically correlated connotations to trust intrinsic to a city: the inner circle represents the area of the highest trust or security, with each of the outer circles corresponding to areas of decreasing trust. This metaphor is then used to categorize users and documents on a distributed file sharing network. According to the metaphor, documents and users encountered in the inner circle of the city are intuitively seen as highly trusted and safe, whereas documents or individuals found in the outer parts of the city can be seen as potentially malicious and should be handled carefully and information regarded skeptically.

We also need to be able to handle documents or users that have undefined trust from the local user. In the metaphor, we do this by introducing the world surrounding the city; this world is partially shrouded in fog and not part of the city hierarchy. Any document or user encountered here is previously unknown to the local user and can potentially fit into any of the circles of the city.

Taking this a step further, we can then use this trust information to classify documents owned by other users in a distributed file sharing system. Using a combination of our own trust assignment policy and the average trust of other users, external documents are classified and added to the same city model. The final result is a dynamic and easily overviewable picture of the trust relationships in the system, giving users a way to easily relate the geographical position of a user or document to its trust level.

Our information visualization technique building on these ideas is called *TrustNeighborhoods*, and is a space-filling radial-layout visualization consisting of concentric rings representing the various trust levels and the buildings on the city grid representing individual documents and users in the system. To simplify the concept of trust, the trust levels are turned into a small number of intervals that we call *societies*. Each society is colored in a such a way to give some indication of the trust level of each society.

The technique can be used to visualize the trust relationships of both users as well as documents in the system. In both cases, the entities are organized into societies according to their trust level. They are then placed in the radial grid of each society ring. Angular placement is controlled by the user to allow for spatial arrangement that makes sense to the user, but could optionally be sorted according to some system property.

Finally, the visualization itself uses this representation of societies consisting of stacked ring arrays to render a graphical image of the system. The *TrustNeighborhoods* method has both a 2D and a 3D mode

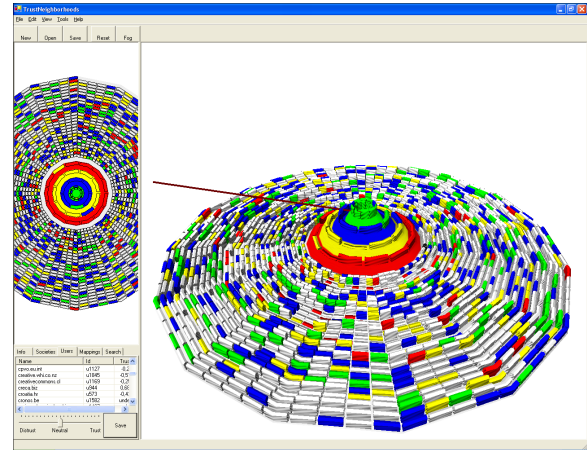


Figure 2: Prototype implementation with 2000 files.

using the same representation; the 2D mode is intended for managing trust and societies, whereas the 3D mode is used for overview and navigation. The behavior of the graphical visualization depends on the mapping between system properties and the graphical actuators; for instance, the user can configure the visualization to map the average trust of a document as the height of the city building representing it. Figure 2 shows an example of our prototype implementation featuring both a 2D and 3D Neighborhoods visualization of 2000 participants in a distributed file sharing system.

### 3.1. 2D Mode: Trust Management

In the 2D mode of the *TrustNeighborhoods* method, we simply draw the societies as concentric rings, using the identifying color of each society as a background. Entities are then drawn in a second pass as circle segments on top of the ring representing the society. If we are visualizing documents, we can draw the circle segment for a document  $D_y$  in the color representing either simply the owner's own trust classification (i.e. we use  $T_x(D_y) = T_y(D_y)$ ), or we can alternatively weigh this value with our own trust in the user (i.e.  $T_x(D_y) = T_x(y)T_y(D_y)$ ). If, on the other hand, we are visualizing users instead of documents, a useful metric for the color of the circle segment is the average trust of the user  $y$  in the system, possibly weighted by our own trust ( $\bar{T}(y)$  or  $\bar{T}_x(y)$ , respectively). The graphical actuators (such as color or height) of the visual object can be mapped to give extra information to the local user, for instance highlighting asymmetric relationships or improper trust assignment ("I trust that user, yet he is not particularly well-trusted among my peers—maybe I should revise my trust?").

Note that societies are disjoint intervals on the trust range  $[-1, +1]$ ; finding the appropriate society and its corresponding color is as simple as finding the interval containing the calculated trust value.

For document trust management, on the other hand, the local user can control the trust levels of the documents she herself owns, but cannot affect the trust level of other documents. She can, however, change the position of the non-local document within the society ring to allow for spatial arrangements that make sense to the user. Changing the trust level of a user will indirectly change the trust level of the documents owned by that user.

Our implementation of the 2D TrustNeighborhoods technique also supports continuous zooming and panning in the visualization to simplify trust management for complex systems. Users can easily pan around and zoom in and out of the visualization to study details.

### 3.2. 3D Mode: Overview and Navigation

The 3D mode of the TrustNeighborhoods technique is primarily designed for overview and navigation. The idea behind this mode is for it to act as a mental aid to a local user accessing documents on the file sharing service; whenever the user is downloading a document, the visualization will pop up to show the context of the document being downloaded and give the user an indication of its trust. In other words, the 3D mode is not intended to serve as a file browser or search interface. Instead, the local user would spend a small amount of initial effort to build up a trust relationship using the 2D mode, and then rely on the system to compute trust appropriately and making it explicit to the user through the 3D mode. Only seldomly would the user have to go back to the 2D mode for explicit trust management. In most cases, the 3D view would be presented in an ambient visual channel, such as the desktop or the window background of the file manager (see Section 6).

The heart of the 3D TrustNeighborhoods visualization is again the stack of societies representing trust levels and the categorization of documents and users into these, but here the metaphor of a city is much stronger than for the 2D mode. The whole data set is rendered on a set of color-coded concentric circles representing the society rings, each ring slightly taller than the one outside it. Entities (i.e. documents or users, depending on which mode the visualization is in) are then rendered as buildings in their respective society rings, mirroring their placement in the 2D version of the visualization. Thus, the user is able to easily recognize the trust level of a specific entity by observing its location in the city, an operation that would

require only a cursory glance at the visualization window.

Each entity has a number of attributes—such as file size, access count, modification date, etc—that can be connected to actuators for the visualization components. These attributes control the graphical look of individual buildings and can thus carry information to the user of varying degrees of subtlety; in our prototype implementation, we use the building color, shape, and height, but more attributes such as texture, size, and details are certainly feasible.

To further aid users in quickly assessing the trust of specific entities, the visualization also supports the use of volumetric fog for the world society. This feature gives a visual indication of the unclear nature of entities in this society, but still affords exact color recognition when the camera draws closer.

### 3.3. Interaction Techniques


In addition to the interactions specified above, the TrustNeighborhoods visualization also supports two additional direct manipulation [Shn83] techniques that are used for interacting with the visualization: a radial distortion mechanism for modifying the visual space assigned to the various society rings (similar to [ET04] and [YWR02]), and a “fly-to” command that is used to smoothly transport users from an overview to a detail view in the 3D mode of the visualization. The purpose of this latter interaction is to provide the user with the context of a specific document being downloaded or a remote user being interacted with in an overview-to-detail fashion.

## 4. User Study

We performed a formal user study of the TrustNeighborhoods technique, comparing subject performance using the tool to their performance when manually assigning trust. We define performance as a combination of time taken to accomplish the task, as well as the correctness of the trust assignment. Users of distributed file sharing systems today typically have no more to go on but the hostname of the remote peer, so this is a reasonable comparison to use. Our hypothesis is that the use of the new technique will improve trust assignment accuracy. We also conducted qualitative testing of the technique, studying how subjects used the tool.

### 4.1. Subjects

We recruited 20 test subjects for this study, two of which were female. The subjects were drawn from the engineering programs of our university and were all in



**Figure 3:** Trust scale with four societies at 0.5 intervals.

the range of 19 to 45 years of age. All subjects had normal or corrected-to-normal vision; one claimed partial color-blindness, but was nevertheless able to successfully complete the test with no reported problems.

#### 4.2. Equipment

The experiment was performed on a dual-processor Intel Pentium 4 Xeon 3 GHz desktop computer with 1024 MB of memory running the Microsoft Windows XP operating system and the TrustNeighborhoods prototype application. The display was a 19-inch LCD flatscreen monitor powered by an NVIDIA Geforce 6600 graphics card; the application was maximized to the resolution of  $1280 \times 1024$ .

#### 4.3. Task

Subjects were asked to assign trust to a sequence of hostnames in the dataset, one for each task. Each task was presented as a short e-mail message where the subject was asked to rank a specific host. The hostnames were collected from the set of hosts with undefined trust from the user. Trust assignment was performed using a continuous slider in the range of  $[-1, +1]$ . A continuous color scale key was also mapped to the slider to give users a visual correlation between the trust scale and the four different trust societies (see Figure 3).

#### 4.4. Dataset

The dataset used for the study was an interesting problem, given that the visualization is based on a hypothetical trust system that does not exist within the context of distributed file sharing systems today. Thus, we were forced to construct our own dataset. To retain validity and repeatability of the experiment, we collected hostnames exclusively from real host databases on the Internet, including mainly so-called DNSBLs (DNS-based Black Lists) for malicious hosts, as well as quality-based ranking sites for the benevolent hosts.

More specifically, malicious hosts were ranked according to their severity (advertisements, spammers, malware, virus sites, and crackers), and hostnames were then collected from the appropriate DNSBLs: the VIRBL (<http://virbl.bit.nl/>) for viruses, bleeding-snort (<http://www.bleedingsnort.com/>), supplying the

Snort IDS, for malware, viruses, and intruders, and the Spamhaus SBL (<http://www.spamhaus.org/sbl/>) for spammers. Benevolent hosts were collected from mini-Rank (<http://minirank.com/>), a popularity-ranking host database based on qualitative factors, taking 50 from each TLD tracked by the site, and ranked according to their popularity.

The distribution of malicious versus benevolent hosts in the combined dataset was another issue; we opted for a 20% versus 80 % distribution for malicious versus benevolent hosts. This figure is loosely based on the average ratio of reports of malicious activity per Internet users reported by the Internet Storm Center (<http://isc.sans.org/>). The full dataset consists of 2000 hosts, randomly selected from the above sources with the given distribution.


In a real scenario, users would initialize the TrustNeighborhoods model with a number of trusted hosts. However, in order to limit the free variables of the user study, the dataset was seeded with 10 fully trusted hosts.

#### 4.5. Design

The experiment was designed as a one-way ANOVA, with the sole independent variable USEVIS (two levels, TRUE or FALSE, respectively), i.e. whether the user had access to the TrustNeighborhoods visualization or not. The dependent variables were the completion time and the trust error, i.e. the absolute difference between the user-assigned trust and the actual trust of a user. Subjects received the USEVIS factor in counter-balanced order to avoid systematic effects of practice.

Each task set consisted of 100 individual tasks of assigning trust to a hostname. The same dataset was used for both conditions, with the hostnames selected randomly from the set so that the same name never occurred twice. For the TrustNeighborhoods condition, the view in the 3D visualization would be smoothly animated to the new hostname, and each trust assignment would lead to the host being added to the corresponding society ring for that trust interval. Furthermore, building height was mapped to average trust and color to weighted average trust. For manual trust assignment, the subject would only have the actual hostname to base the decision on. Each task was individually timed, starting from when the new hostname was given until the user had performed the trust classification.

Each session lasted approximately 30 to 45 minutes. Subjects were given a training phase of up to 10 minutes during which they were instructed in the use of the visualization as well as given a test run of the trust assignment task for both conditions. During the actual



**Figure 4:** Average error and average completion times; error bars show standard deviations.

test, only general questions on the visualization were answered by the test administrator, and no assistance on specific tasks was given.

With 20 participants and 100 trust assignments for each of the two levels of the USEVIS factor, 4000 individual tasks were recorded in total. Upon completing both task sets, subjects were asked to fill out a post-test questionnaire.

#### 4.6. Qualitative Tests

Due to the lack of an obvious existing tool or technique to use as a comparison condition for the formal user study, we also endeavoured to observe all participants closely during each test session as well as interviewing them briefly upon completing the study. The purpose of this was to identify potential problems in the representation as well as to give an additional measure of how subjects used the tool and how they perceived its benefits and drawbacks.

### 5. Results

Our results show that subjects were significantly more correct in their trust assignment when using the new technique than for manual assignment, but at the cost of additional time investment.

#### 5.1. Correctness

The average error for both conditions was 0.45 (s.d. 0.17). Using manual assignment (i.e. with USEVIS = FALSE), the average error was 0.57 (s.d. 0.09), whereas the average error was 0.33 (s.d. 0.13) when subjects used the visualization (USEVIS = TRUE). This is a significant difference;  $F(1, 19) = 80.98, p < .001$ .

#### 5.2. Completion Times

The average completion time for a single trust assignment task was 6.92 (s.d. 3.29) seconds. For the manual condition using no visualization, the average completion time was 4.84 (s.d. 2.00) seconds, versus 9.24 (s.d. 2.76) seconds when using the visualization. This is a significant difference;  $F(1, 19) = 91.41, p < .001$ .

### 5.3. Subjective Ratings

Table 1 summarizes the subject rankings from the post-test questionnaire; all were significant to  $p < .05$ . In summary, the new technique was the most preferred method in all aspects.

Attribute	Manual	TrustN
Q1. Rank w.r.t. ease of use	0.25	0.75
Q2. Rank w.r.t. efficiency	0.05	0.95
Q3. Rank w.r.t. enjoyment	0.10	0.90
Q4. Rank w.r.t. speed	0.75	0.25
Q5. Rank overall	0.10	0.90

**Table 1:** Subjective rankings.

### 5.4. Qualitative Results

Subjects had no problem understanding the general idea of the technique, and were quick to grasp the concepts of color coding, entity size, and the society ring structure (in all cases well within the assigned ten-minute training phase). None of the involved subjects thought that the metaphor felt unnatural. However, many users seemed to have problems with the 3D navigation controls in the application, even though the controls have been constrained to simple orbit and zoom operations; it is clear that they must be constrained further (or eliminated completely) in a real application.


Some subjects expressed frustration over the visual representation and the interaction, naming it unwieldy and slow compared to the manual approach. These were generally especially computer-savvy subjects with a broad computer background. A tendency among this group of subjects was also a general skepticism towards the visualization; instead of accepting the indications given by the visualization for a particular trial, the subject would go with his or her own reasoning (mostly in a less trusting manner).

Overall, subjects seemed happy about the visual representation, and many remarked that they found it odd that nothing similar was available for general use. A few expressed an interest in using a TrustNeighborhoods-based tool for their daily work.

### 6. User Environment Integration

Our vision for the use of TrustNeighborhoods is to integrate it with existing user desktop environments such as Microsoft Windows and MacOS. The user would be able to assign trust to just a few known users with the help of a standard interface for the distributed file sharing mechanism. Whenever the user requests access to a remote file on the network, the





**Figure 5:** Mockup of a TrustNeighborhood view running in the background of a UNIX graphical file explorer. (Note that this is a composite screenshot.)

“fly-to” interaction of the 3D view could be invoked in the background of the desktop environment, non-obtrusively providing the user with the trust context of the file.

Our prototype implementation runs as a stand-alone application and has not yet been integrated into a real user environment, but we have made some prototype mockup screenshots of our ideas. Figure 5 shows one such from a Linux system, with the TrustNeighborhoods view integrated in the background of the Nautilus file explorer itself. Here, traditional file browsing is augmented with the 3D TrustNeighborhoods view showing the context and appearance of a particular file being downloaded from a public FTP server. Again, this appears to be a safe file the user can download with no extra precautions needed.

## 7. Discussion

When attempting to visualize an abstract concept such as trust, the choice of visual mapping falls entirely on the researcher. Our method for designing the mapping has in this project been to employ metaphors to guide the process, allowing especially novice users to transfer existing knowledge to this new domain. For this reason and from the results and comments collected from the participants of the study, we believe the chosen mapping to be sound, but other visual mappings are certainly possible.

Subjects using the TrustNeighborhoods technique were almost twice as accurate in their trust assignments than when manually assigning trust based on the hostname alone. This comes at a certain time tradeoff, however. This is yet another example of the classic time versus quality tradeoff; in general, the

more time is put into a task, the more accurate the result will be. Our technique helps users make more informed decisions, but this extra information means they will take longer to do so. On the other hand, the cost for mistakes can be costly, potentially leading to loss or corruption of the personal data of the user, and thus this may be an acceptable tradeoff.

Regarding the construction of the dataset used in the study, it is clear that it would be better to use a real dataset constructed for the purpose. However, it is important to keep in mind that the visualization technique itself is more or less decoupled from the dataset and the underlying trust model. Hence, we expect similar results would arise from different choices of trust model and dataset.

Our visualization is aimed primarily at novice and intermediate users, yet we made use of engineering students with an overall high computer background in the study. It also meant that the study involved a few “power users” who were well-versed in computers and Internet security, and to whom a visual representation is almost an encumbrance (as well as a potential source of suspicion—an informal observation is that the more experienced the user, the less trusting he or she is). This may have an impact on the generality of our results. Nevertheless, it is our belief that the outcome of an experiment with only novice users would be similar to the current one.

Our qualitative evaluation indicated that some subjects were very skeptical of the visualization and some employed a method of almost doing the reverse of what the tool told them—this was particularly true of users with very high computer experience. This is an interesting and unexpected effect that deserves closer scrutiny in future studies. It may be that experienced users are especially wary of deceptive information and have a dislike of being “told what to do”, and thus are suspicious of any tool that attempts to do so.

## 8. Conclusions and Future Work

We have presented TrustNeighborhoods, an information visualization technique based on circles of relationships implemented using a city metaphor. This is done in order to show the trust level of entities in a distributed file sharing system by relating them to their geographical position in the city. The technique is aimed at novice and intermediate-level users, helping them acquire a mental model of the trust and surrounding context of documents they are accessing on the network. We have implemented a prototype of the technique, and our user study indicates that the method is a significant improvement over manual trust assignment. To summarize, the main contributions of this paper are the following:

- an adaptation of Shneiderman's circles of relationship to a tangible city metaphor useful for information visualization;
- an intuitive and flexible technique for trust visualization building on this city metaphor; and
- a formal user study showing that trust assignment using the visualization is significantly more correct than for manual trust assignment, albeit at the cost of additional time investment.

## Acknowledgments

Thanks to the subjects who participated in the user study and offered valuable insight on how to improve the visualization. Special thanks to James Eagan of Georgia Tech for his thorough review of the paper and his suggestions for improving the work.

## References

- [AH98] ANDREWS K., HEIDEGGER H.: Information slices: Visualising and exploring large hierarchies using cascading, semi-circular discs. In *Proceedings of the IEEE Symposium on Information Visualization 1998 (Late-Breaking Hot Topics)* (1998), pp. 9–12.
- [Chu98] CHUAH M.: Dynamic aggregation with circular visual designs. In *Proceedings of the IEEE Symposium on Information Visualization 1998* (1998), pp. 35–43.
- [ET04] ELMQVIST N., TSIGAS P.: *CiteWiz: A Tool for the Visualization of Scientific Citation Networks*. Tech. Rep. CS:2004-05, Chalmers University of Technology, 2004.
- [Fur86] FURNAS G. W.: Generalized fisheye views. In *Proceedings of the ACM CHI'86 Conference on Human Factors in Computer Systems* (1986), pp. 16–23.
- [GK03] GOOD N. S., KREKELBERG A.: Usability and privacy: a study of Kazaa P2P file-sharing. In *Proceedings of ACM CHI 2003 Conference on Human Factors in Computing Systems* (2003), pp. 137–144.
- [KvdWvW01] KLEIBERG E., VAN DE WETERING H., VAN WIJK J. J.: Botanical visualization of huge hierarchies. In *Proceedings of the IEEE Symposium on Information Visualization 2001* (2001), pp. 87–96.
- [LR96] LAMPING J., RAO R.: The hyperbolic browser: A focus + context technique for visualizing large hierarchies. *Journal of Visual Languages and Computing* 7, 1 (1996), 33–35.
- [LRP95] LAMPING J., RAO R., PIROLI P.: A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of the ACM CHI'95 Conference on Human Factors in Computing Systems* (1995).
- [Mar94] MARSH S. P.: *Formalising Trust as a Computational Concept*. PhD thesis, Department of Computer Science and Mathematics, University of Stirling, 1994.
- [MB95] MUNZNER T., BURCHARD P.: Visualizing the structure of the World Wide Web in 3D hyperbolic space. In *Proceedings of the Symposium on Virtual Reality Modeling Language 1995* (1995), pp. 33–38.
- [RMC91] ROBERTSON G. G., MACKINLAY J. D., CARD S. K.: Cone trees: Animated 3D visualizations of hierarchical information. In *Proceedings of the ACM CHI'91 Conference on Human Factors in Computing Systems* (1991), pp. 189–194.
- [SCGM00] STASKO J., CATRAMBONE R., GUZDIAL M., MCDONALD K.: An evaluation of space-filling information visualizations for depicting hierarchical structures. *International Journal of Human-Computer Studies* 53, 5 (2000), 663–694.
- [Shn83] SHNEIDERMAN B.: Direct manipulation: A step beyond programming languages. *IEEE Computer* 16, 8 (1983), 57–69.
- [Shn02] SHNEIDERMAN B.: *Leonardo's Laptop: human needs and the new computing technologies*. MIT Press, 2002.
- [sig05] sig2dot, 2005. <http://www.chaosreigns.com/code/sig2dot/>.
- [SZ00] STASKO J., ZHANG E.: Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Proceedings of the IEEE Symposium on Information Visualization 2000* (2000), pp. 57–68.
- [TK02] TAKADA T., KOIKE H.: Tudumi: Information visualization system for monitoring and auditing computer logs. In *Proceedings of the International Conference on Information Visualization* (2002), pp. 570–576.
- [WT99] WHITTEN A., TYGAR J. D.: Why Johnny can't encrypt; A usability evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium* (1999), pp. 169–184.
- [Yee02] YEE K.-P.: User interaction design for secure systems. In *International Conference on Information and Communications Security (ICIS)* (2002).
- [YWR02] YANG J., WARD M. O., RUNDENSTEINER E. A.: InterRing: An interactive tool for visually navigating and manipulating hierarchical structures. In *Proceedings of the IEEE Symposium of Information Visualization 2002* (2002), pp. 77–84.