

# Grand Central Usenet: The Design and Evaluation of a Thread-Based Usenet Browser

**Carman Neustaedter**  
University of Calgary  
2500 University Drive NW  
Calgary, AB, Canada T2N 1N4  
carman@cpsc.ucalgary.ca  
+1 403-210-9501

**Marc A. Smith, Gina Danielle Venolia**  
Microsoft Research  
One Microsoft Way  
Redmond, WA 98052  
{masmith, ginav}@microsoft.com  
425-706-8080

## ABSTRACT

Interfaces to online discussion spaces, such as email discussions, lists, and newsgroups, do a poor job of representing the structure and temporal development of conversation threads. These limitations contribute to user overload and to the erosion of the value of these channels. In this paper, we present an alternative interface to threaded conversations, Grand Central Usenet, which features a graphical interface component that highlights the size, structure, and development of conversation threads. We harnessed this interface to Usenet newsgroup data and conducted a user study that contrasted this interface with a standard message browsing tool. Users showed significant improvements in productivity, reports of ease-of-use, and satisfaction with our design in contrast to a widely used standard interface.

## Author Keywords

Threaded conversations, email, web boards, Usenet newsgroups, visualization, user interface, social computing

## INTRODUCTION

Online discussions, whether implemented as email, email lists, web boards, or newsgroups, can unfold over long periods of time with a shifting population of participants. However, interfaces to such data structures often lack any indication of the central attributes of threaded conversations; their size, structure, and population. As a result, threads are often presented in fragmented ways that cause confusion and limit the utility of the channel for complex discussions.

An alternative approach focuses explicitly on the structure of threaded conversations, highlighting the turn and reply patterns and their unfolding over time. Prior work by

Venolia and Neustaedter [13] presented a thread visualization integrated into an email based conversation application. In our current work, we report efforts to adapt and evaluate this interface to the related data structures found in Usenet newsgroups. While email lists and newsgroup conversations share many properties, this involved some adaptations to the unique qualities of Usenet. Our design is a response to the problem posed by many conversation repositories like Usenet, in which collections of threads are so large that it is difficult to select information likely to be interesting.

Our approach, called Grand Central Usenet, brings together several key features that augment the standard tools available in existing message based browsers. Specifically, we combine an effective visualization of thread structures with social accounting metadata which describe aggregate properties of the conversational environment; for example, the number of messages in a discussion space or the number of messages an individual author contributes to a collection of threads, are measures of various dimensions of a conversational social cyberspace.

To assess the design of Grand Central Usenet, we created and deployed a user study to capture the differences in the productivity and satisfaction users had with this tool in contrast with a commonly available message browser, Microsoft Outlook Express. The results of our study show that Grand Central Usenet provides users with increased productivity, ease-of-use, and user satisfaction

In the following, we describe the Grand Central Usenet interface, the user study we designed and implemented to evaluate the interface, and our findings which show significant improvements in user experience.

## RELATED WORK

Related work has also focused on alternative methods for displaying richer information about threads and conversational spaces.

The original Loom project [1, 2] focused on visualizing social patterns within Usenet newsgroups. It highlighted saliencies such as rowdy, vociferous users as well as the number of participants in different threads over time. It also visualized the difference between initiated posts and replies.

Loom managed to uncover interesting author dynamics found in newsgroups – for instance, the marked difference between the average number of participants per thread in technical versus social newsgroups.

Conversation Map [9] also looks at Usenet newsgroups and touches on the issue of the people present in the conversational space. It computes a set of social networks detailing who is talking to whom and who is citing whom in the newsgroup. The other main feature in Conversation Map is its visualization of the centrality degree of users in the newsgroup where the social network of each newsgroup allows us to understand which users are more central than others to that group’s discussions. Here, as in Loom, remarkable patterns emerge that are related to people’s interactions in the conversational space.

PeopleGarden [15] visualizes message boards in terms of their authors’ activities. Each flower in PeopleGarden represents a user in the conversational space and its petals represent his/her postings. PeopleGarden also shows the amount of replies to a user’s post by displaying pistil-like circles on top of a petal to denote responses. Even though PeopleGarden’s focus was not Usenet newsgroups, the fact

that it strived to identify a conversational space in terms of its people makes it conceptually close to the work presented here.

The Netscan project [12] data mines Usenet and parses each message’s header into a SQL database. Aggregations are performed to produce social accounting metrics about authors, newsgroups, and threads. The Netscan dashboard produces several thread visualizations including an interactive thread tree, a piano roll showing author participation in threads, and a sociogram showing author relationships.

Additional work visualizing online discussions has also been done by [4, 5, 6, 8, 11, 14].

### USER INTERFACE DESIGN

Grand Central Usenet features several user interface components designed specifically to support thread-based browsing of messages, e.g., email, instant messages, and newsgroups posts. Grand Central Usenet was previously described by Venolia and Neustaedter [13]. Since then, we have adapted the user interface and backend of the system to support visualizing threads from Usenet newsgroups. The backend of Grand Central Usenet now connects to

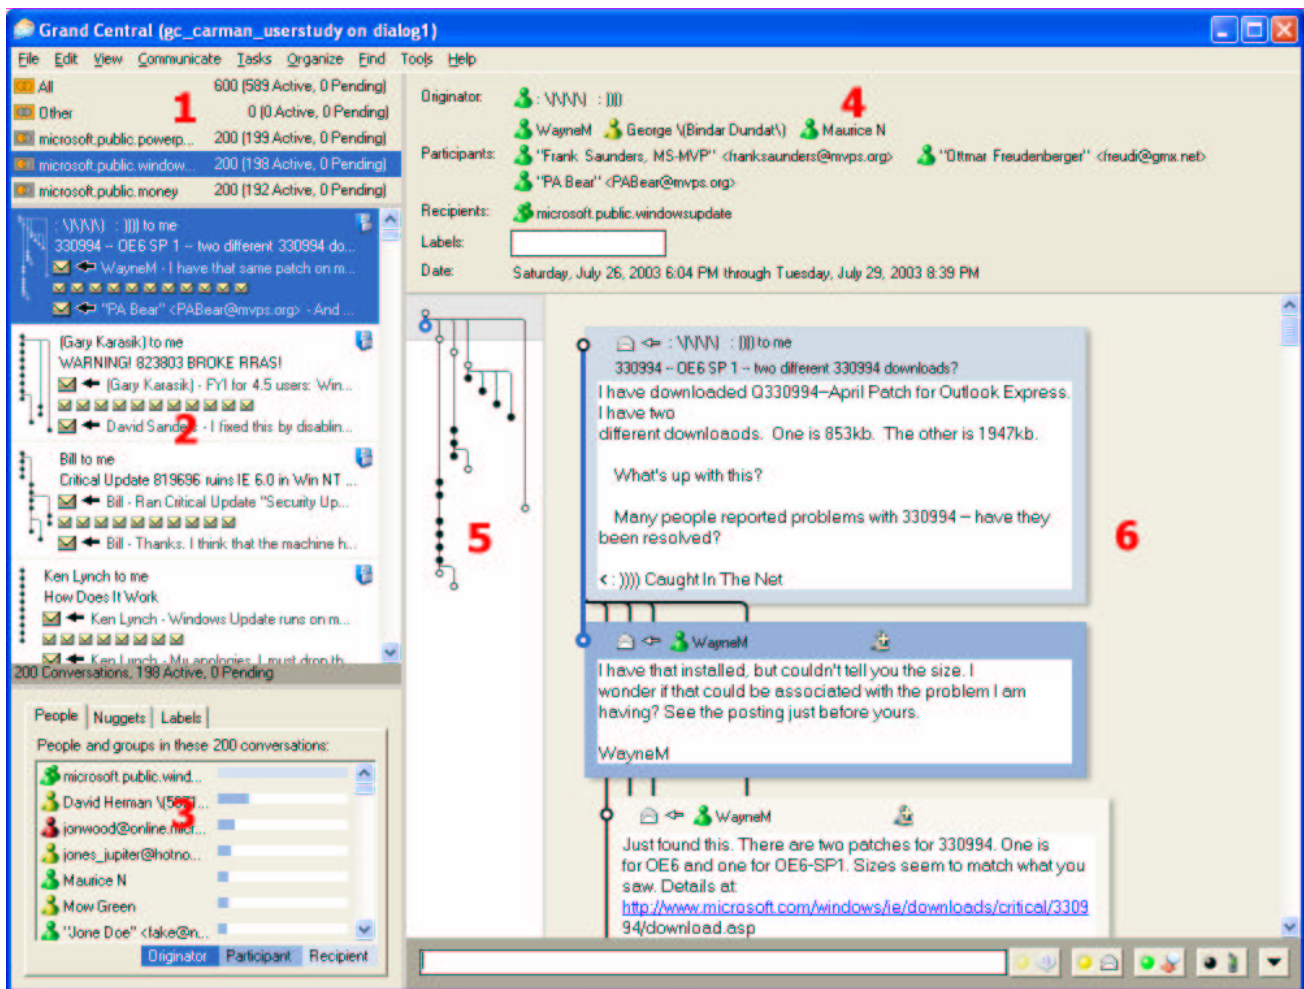


Figure 1: Grand Central Usenet showing newsgroup threads. The numbered sections are described in the text.

Netscan's server to retrieve newsgroup data and social accounting metrics. Figure 1 shows the main Grand Central Usenet interface and Figure 2 shows a secondary window used for subscribing to newsgroups.

### Newsgroup Selection

Users select newsgroups to view using a secondary window shown in Figure 2. The left pane of this window contains a set of dynamic query filters [10] where users can filter newsgroups based on criteria such as keywords in the newsgroup name or social accounting metrics. The query is dynamic in the sense that results are automatically updated as users change the search criteria. This provides a tight-coupling between query parameters and search results, and allows for easy search refinement. Currently, the interface allows users to filter based on two metrics, the number of posts and the number of posters or authors, yet more could easily be added.

Newsgroups that match the search criteria are displayed in the table on the right pane of the window along with associated social accounting metrics (Figure 2, the right pane from left to right): the number of posts, posters or authors, the ratio of posters-to-posts, the number of people

posting in the current time period who also posted in the previous time period, the number of replies, and people who have replied. Clicking on any column heading causes the results to sort according to the values for that heading.

A table lens [7] allows users to view the search results at varying levels of detail. Figure 2 shows the table fully collapsed with several rows expanded by the user. Fully expanded results show complete newsgroup names along with numeric values for each social accounting metric. Results can be collapsed in the user interface using controls in the bottom left corner of Figure 2. Fully collapsed results hide newsgroup names and show each social accounting metric as a representative colored bar. The fully collapsed view allows for visual comparisons and an increase in the amount of viewable results at any one time. A slider allows users to select the collapsed size, which will alter the height of each row, thus increasing or decreasing the amount of results viewable at any one time. If the size of each row is large enough, newsgroup names will be visible. Individual rows in the table can be expanded or collapsed by simply clicking on the given row (Figure 2).

Users can choose to subscribe to a newsgroup by clicking



Figure 2: Grand Central Usenet's subscription interface features a dynamic query of newsgroups based on social accounting metrics and a table lens [7] for visualizing the search results.

on a newsgroup name within the results table. A similar action will unsubscribe the user from it.

### Newsgroup and Thread Browsing

Newsgroup and thread browsing is accomplished using the main Grand Central Usenet window, shown in Figure 1. Containers are created automatically to hold threads for each subscribed newsgroup, (1) in Figure 1. Users are also able to create new containers and place threads in multiple containers. For example, a thread from the “*rec.sports.soccer*” newsgroup could also be placed in a user-created “*Hobbies*” container. Threads for a selected container are shown in (2) with a thumbnail of the thread’s tree, the subject line, the name of the thread originator, and information about unread messages in the thread. Social accounting data about authors for a given container are shown in (3) where users can see how many threads an author has started and the number of replies the author has posted. These metrics can help determine the quality of each message’s content [3]. Filtering and sorting based on author metrics is provided in context-menus.

Selected threads appear on the right side of Figure 1 with header information about the thread shown in (4), e.g., the originator, participants, and time span. Messages for the selected thread are shown in the thread visualization in (6). Venolia and Neustaedter describe this component in detail and evaluate it in terms of its legibility in [13]; messages are embedded within a mixed-model visualization of the thread tree showing both the reply-structure and temporal nature of the thread (the list is chronological with newest messages appearing at the bottom). A smaller interactive thread tree is shown in (5): users can select nodes within (5), causing (6) to scroll to the selected message. A grey band in (5) shows which messages are currently visible in (6). This band also acts as an in-context scroll bar where dragging it causes (6) to scroll.

### USER STUDY

To test whether people can perform newsgroup and thread tasks better with Grand Central Usenet, we designed and implemented a user study to compare it to Microsoft Outlook Express (OE). OE is a widely used newsgroup browser, which provides a good example of the features commonly found in most other newsgroup browsing tools. Our study compares the main interfaces of Grand Central Usenet and OE, but does not evaluate and contrast the newsgroup selection interfaces given the lack of an analogous newsgroup selection interface in the OE interface.

### Methodology

Participants performed a series of tasks with both OE and Grand Central Usenet. We tested several dependent variables per task for each browser: time for completion, accuracy, and the number of unique threads and messages viewed. We also tested two additional dependent variables, as rated by each user: ease of use, and user satisfaction.

### Materials: User Tasks

Users performed four realistic newsgroup tasks based on user feedback from a technical support team at our company and our own personal experience and background knowledge of Usenet. We included two types of tasks: high level tasks (Task 1 and 2) where the user’s goal was to gain a higher level understanding of a newsgroup, and low level tasks where the user’s goal was to gain a detailed understanding of a thread, be it small (Task 3) or large (Task 4). Users were asked to:

1. *Hot Topics*: Find three “hot topics” in this newsgroup, e.g., the three largest threads (those with the most messages).
2. *Top Author*: Find the author who has posted the most messages in this newsgroup.
3. *Small Thread*: Find the asked for thread (described below) in this newsgroup. How many messages does this thread have? How many different people posted messages in this thread? Has the conversational topic changed? Do any conversational topics remain unresolved?
4. *Large Thread*: Find the largest thread in this newsgroup and answer the same questions about this thread as the previous task.

### Materials: Usenet Data

One of the most difficult aspects of designing the study was finding comparable data sets for participants to view with each newsgroup browser. Both sets of data needed to have a similar overall number of threads and messages, along with threads of a similar branching nature and size, and a variety of thread types, e.g., small, large, bushy, linear. We also wanted to select data that was general enough that a large majority of participants would understand it easily. This meant we had to select data from newsgroups discussing topics familiar to most, e.g., Microsoft Word, Windows XP. We carefully selected Usenet data from two newsgroups for users to view for each browser.

For set A, we preloaded data from the microsoft.public.money newsgroup for the period of July 24 – 28 into both OE and Grand Central Usenet. In OE, this meant downloading a set of *messages* with a received date in this time period. Since Grand Central Usenet is a thread-based browser, this meant downloading a set of *threads* with a start date that fell in this time period. While both browsers did not display identical data for the same newsgroup, due to the nature of each browser be it message-based or thread-based, the data shown in both browsers was arguably comparable. One notable difference was that OE typically contained broken threads where messages that had a received date outside of the selected time period were not included. In contrast, Grand Central Usenet contained only complete threads. Data set A contained approximately 200 threads and 500 messages in both OE and Grand Central Usenet.

In Grand Central Usenet, the largest thread in set A contained 29 messages posted by 15 different authors. In OE, the same thread was the largest, although it appeared broken with only 17 messages posted by 10 different authors. We normalize our results to account for these differences. For Task 3, *Small Thread*, we selected a thread with 5 messages posted by 3 different authors, which contained two branches. This thread appeared identical in both Grand Central Usenet and OE.

For set B, we preloaded data from the microsoft.public.windowupdate newsgroup for the period of July 24 – 28 into both OE and Grand Central Usenet. Set B contained approximately 200 threads and 550 messages in both OE and Grand Central Usenet.

In Grand Central Usenet, the largest thread in set B contained 22 messages posted by 7 different authors. In OE, the same thread was the largest, yet it appeared broken and only contained 11 messages posted by 4 different authors. We normalize our results to account for these differences. For Task 3, *Small Thread*, we selected a thread with 4 posts, 3 of which were all replies to the original message. This thread appeared identical in both Grand Central Usenet and OE.

The microsoft.public.powerpoint newsgroup was used for a training session with each browser during the study. We preloaded both browsers with data from this newsgroup for the same time period of July 24 – 28. The training data set contained approximately 200 threads and 650 messages in both browsers

#### **Materials: Questionnaires**

A pre-test questionnaire gathered demographics about each participant, as well as background information such as how frequently they browsed newsgroups, what their favorite newsgroup browser was, and how experienced they were with OE.

A post-browser questionnaire asked participants for their subjective rating of a browser's ease-of-use and also assessed user satisfaction with the browser. The post-browser questionnaire contained a series of statements, e.g., "I found it easy to use this browser," where users rated their agreement on a 5-point Likert scale (1-strongly disagree to 5-strongly agree).

The post-test questionnaire asked participants to choose which browser they preferred for various scenarios, e.g., viewing simple threads, viewing complex threads, gaining an overview of a newsgroup. Participants could choose OE, Grand Central Usenet, or neither browser.

#### **Method**

Prior to arriving at the user study, participants completed the pre-test questionnaire online. We then randomly assigned each participant to one of four groups to counterbalance which browser participants would use first and which set of data they would view with each browser. In order to make these two browsers as comparable as

possible, we set the default viewing mode in OE to "Group by Conversation" and defaulted to auto-expand each conversation group.

First, participants performed a short training session with one of the two browsers where the browser's user interface and main functionality was explained. They were then given the opportunity to explore the training newsgroup with the current browser. Second, participants were given each of the four tasks in turn to perform. Third, participants answered the post-browser questionnaire about the browser they had just used. These steps were then repeated for the second browser. The study concluded after the completion of the post-test questionnaire.

#### **Participants**

We recruited 16 participants, 9 male and 7 female, from within our company through an email announcement widely distributed through internal email lists. Participants completed a preliminary survey and were selected based on their availability and experience with newsgroups. Participants held a variety of occupations including software development, management, and technical support. Participants were all regular Usenet users most frequently browsing newsgroups for work-related purposes, while occasionally for non-work related purposes. All were familiar with the concepts of newsgroup threads and thread trees.

We validated the practicality of our user study tasks by asking users to list the typical tasks that they perform with Usenet in the pre-test questionnaire.

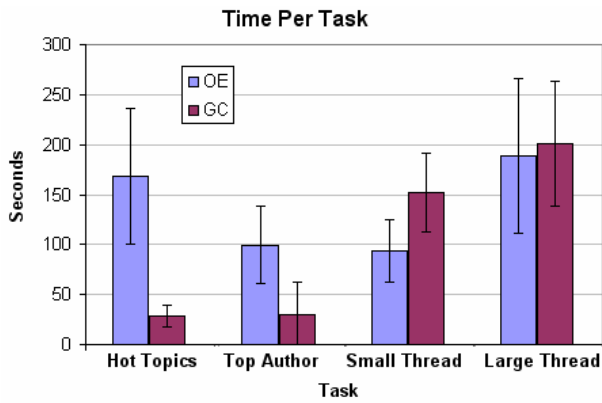
#### **RESULTS**

Our results show that Grand Central Usenet provides increased user productivity, satisfaction, and ease-of-use for most of the functionality we tested in our user study. We now discuss each of our findings in detail.

#### **Time Per Task**

Figure 3 shows the mean times needed to complete each task. A series of paired-samples T-tests shows participants were able to complete Task 1, *Hot Topics*, ( $p < 0.05$ ,  $DF=15$ ,  $t=8.02$ ) and Task 2, *Top Author*, ( $p < 0.05$ ,  $DF=15$ ,  $t=4.90$ ) faster with Grand Central Usenet than OE and the differences are statistically significant. Participants were faster at completing Task 3, *Small Thread*, ( $p < 0.05$ ,  $DF=15$ ,  $t=-5.69$ ) with OE than Grand Central Usenet and again the results were significant. No significant differences were found in the times for Task 4, *Large Thread*, ( $p = 0.550$ ,  $DF=15$ ,  $t=-0.611$ ).

*Large Thread* asked participants to answer questions about the largest thread and this task required most participants to read almost all messages in each thread. As noted, threads can appear broken in OE because it is a message-based browser. This is not the case for Grand Central Usenet and as a result the threads analyzed by users for *Large Thread* were larger in Grand Central Usenet. For this reason, we normalized the times for this task by calculating the time/message taken by users. The mean time/message for



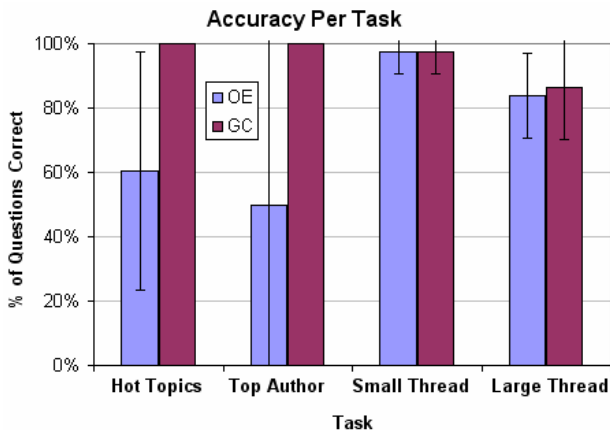
**Figure 3:** The mean and standard deviations of the participant times for completing each task (n = 16).

Grand Central Usenet was  $8.1 \pm 2.6$  s, while OE was  $13.3 \pm 4.2$  s. A comparison of these means shows that users performed *Large Thread* significantly faster ( $p < 0.05$ ,  $DF=15$ ,  $t=5.67$ ) with Grand Central Usenet than OE on a “per message” basis.


### Accuracy

*Hot Topics* was scored out of three based on the correct naming of the three largest threads. *Top Author* was scored out of one based on the correct name of an author. *Small Thread* was scored out of five: one point for finding the asked about thread, one point for the correct number of messages, one for the correct number of authors, one for a sufficient yes or no answer about topic drift, and one for a sufficient yes or no answer about unresolved topics. *Large Thread* was scored the same as *Small Thread*. Participants were told they didn’t have to answer a question if they gave up, or refused to because the task was too onerous; participants received zero points on the subtask if they chose not to answer.

Figure 4 shows participants’ mean accuracy for each of the four tasks. Participants were significantly more accurate when using Grand Central Usenet than OE for *Hot Topics*



**Figure 4:** The mean and standard deviations of participant accuracy for each task (n=16).



**Figure 5:** The mean and standard deviations of the number of unique threads viewed by participants for each task (n = 16).


( $p < 0.05$ ,  $DF=15$ ,  $t=-4.29$ ) and *Top Author* ( $p < 0.05$ ,  $DF=15$ ,  $t=-3.87$ ). No significant differences were found between browsers for *Small Thread* ( $p = 1$ ,  $DF=15$ ,  $t=0$ ) and *Large Thread* ( $p = 0.61$ ,  $DF=15$ ,  $t=-0.522$ ).

Three participants “gave up” on *Hot Topics* with OE and were not able to provide answers for the task. Three participants (although not all the same) also “gave up” on *Top Author* with OE and were not able to provide an answer. No participants gave up when using Grand Central Usenet for any of the tasks.

### Threads and Messages Viewed

Figures 5 and 6 show the mean number of unique threads and messages, respectively, that users viewed to complete each task. Participants viewed significantly fewer threads with Grand Central Usenet than OE for *Hot Topics* ( $p < 0.05$ ,  $DF=15$ ,  $t=2.47$ ), *Small Thread* ( $p < 0.05$ ,  $DF=15$ ,  $t=3.084$ ), and *Large Thread* ( $p < 0.05$ ,  $DF=15$ ,  $t=5.337$ ). No significant difference was found in the number of threads viewed for *Top Author* ( $p = 0.15$ ,  $DF=15$ ,  $t=1.50$ ).

Participants viewed significantly fewer messages with Grand Central Usenet than OE for *Hot Topics* ( $p < 0.05$ ,  $DF=15$ ,  $t=2.218$ ) and *Small Thread* ( $p < 0.05$ ,  $DF=15$ ,



**Figure 6:** The mean and standard deviations of the number of unique messages viewed by participants for each task (n = 16).

$t=3.15$ ). No significant differences were found for *Top Author* ( $p = 0.15$ ,  $DF=15$ ,  $t=1.49$ ) and *Large Thread* ( $p = 0.815$ ,  $DF=15$ ,  $t=-0.238$ ).

Again, threads analyzed for *Large Thread* were larger in Grand Central Usenet than OE; thus, participants had the opportunity to view more messages in Grand Central Usenet than OE. To normalize this, we compared the percent of the total messages for the thread which were viewed in each browser. In OE, participants, on average, viewed  $98.3 \pm 32.7$  % of the messages in the largest thread to complete the task. In Grand Central Usenet, participants, on average, viewed  $57.5 \pm 31.5$  % of the messages in the largest thread. The difference between these means is statistically significant ( $p < 0.05$ ,  $DF=15$ ,  $t=4.78$ ); participants viewed a smaller percent of the messages when using Grand Central Usenet to complete *Large Thread*.

### Subjective User Rating

We asked participants to assess how easy it was to perform aspects of each task with the browser. Table 1 shows the means responses for a selected set of statements found in the post-browser questionnaire, answered on a 5-point Likert scale (1-strongly disagree to 5-strongly agree).

We used a series of Wilcoxon-signed rank tests to compare the responses given for each browser. We found no significant difference in subjects' responses about how easy Grand Central Usenet was to learn compared to OE ( $p=0.102$ ,  $z=-1.633$ ,  $n=16$ ). As well, no significant differences were found for ease of use between browsers ( $p=0.119$ ,  $z=-1.560$ ,  $n=16$ ). Grand Central Usenet made it easier than OE for participants to find large threads ( $p < 0.05$ ,  $z=-3.568$ ), authors with the most posts ( $p < 0.05$ ,  $z=-3.342$ ), and unreplied messages ( $p < 0.05$ ,  $z=-2.661$ ). As well, it was also easier to get an overview of each thread's structure and understand its branching ( $p < 0.05$ ,  $z=-2.862$ ).

When asked what people liked about Grand Central Usenet

Statement	Outlook Express	Grand Central Usenet
It was easy to learn to use.	$3.81 \pm .544$	$4.06 \pm .250$
It was easy to find large threads.	$1.69 \pm .479$	$4.75 \pm .447$
It was easy to find authors with the most posts.	$1.94 \pm .929$	$4.56 \pm .892$
It was easy to find messages without replies.	$2.69 \pm .946$	$3.94 \pm 1.124$
It was easy to get an overview of a thread's structure.	$3.19 \pm 1.047$	$4.50 \pm .516$

**Table 1:** Participant mean ease-of-use responses for both browsers (1-strongly disagree to 5-strongly agree,  $n = 16$ ).

they listed: easy navigation of threads, the thread visualizations, and the use of thread-based metrics. Dislikes included performance issues when scrolling, minor software bugs, and the current lack of text search features.

We also asked participants if they would prefer each browser over their current newsgroup browser. For regular users of OE, 4 out of 6 preferred Grand Central Usenet. All 3 users of a company-internal newsgroup tool preferred Grand Central Usenet. One user of Google Groups preferred Grand Central Usenet, while the other was undecided. The two users of Microsoft WebNews and the one user of XNews preferred Grand Central Usenet. Nobody preferred OE over their current newsgroup browser.

We also asked participants to choose which browser they preferred, given a certain task situation. Table 2 shows these results. More participants preferred Grand Central Usenet for gaining an overall understanding of a newsgroup and thread, and viewing and understanding complex threads. More participants, albeit only one, preferred OE for viewing and understanding simple threads. Overall, participants overwhelmingly preferred Grand Central Usenet because they felt it offered easier navigation, better sorting options, and was more intuitive to use. Participants who preferred OE thought that OE's performance was better (e.g., speed, reliability) and that Grand Central Usenet's user interface was somewhat overloaded.

### DISCUSSION

The results of our user study clearly show that Grand Central Usenet provided users with increased productivity, ease-of-use, and user satisfaction compared to OE. We believe that this is largely due to the data being organized in threads rather than messages. This increase is primarily found when users deal with large or complex threads, or attempt to gain a high-level understanding of a newsgroup or thread.

First, we found that participants could perform most tasks faster with Grand Central Usenet and it required them to view fewer messages and threads. This performance

Task	Outlook Express	Grand Central Usenet	Neither
Gaining an overview of a newsgroup.	1	14	1
Gaining an overview of a thread.	2	13	1
Viewing and understanding simple	7	6	3
Viewing and understanding complex threads.	3	13	0
Overall	2	13	1

**Table 2:** The number of participants who preferred each browser for a given task ( $n = 16$ ).

increase was primarily found when viewing large or complex threads. In the case of simple threads, participants could perform tasks faster with OE, yet participants had to view more threads and messages with OE than Grand Central Usenet. Faced with a deluge of information it is necessary to read less chaff and more wheat. Using thread structure as an additional cue, users are able to select fewer messages and still get the sense of the conversation flow by selecting the most structurally significant messages.

Participants also saw an increase in accuracy when using Grand Central Usenet for performing high-level tasks. No improvements were found for low-level tasks. Participants found it easier to use Grand Central Usenet than OE for performing all of the task scenarios that we gave them. Moreover, most participants preferred Grand Central Usenet over their current browser, be it OE, Google Groups, or another newsgroup browser. When comparing OE to Grand Central Usenet, almost all participants preferred Grand Central Usenet to perform the tasks tested in the user study.

#### **FUTURE WORK**

The results suggest that our design has potential as an alternative interface to complexly threaded conversation environments. Our future work involves componentizing each feature of our main user interface so that future thread-based designs may more easily incorporate many of the ideas we have presented. We also plan to continue our development of Grand Central Usenet by giving users the ability to post messages, integrating additional social accounting metrics into the user interface, and integrating related newsgroup visualizations of author reputation and other newsgroup attributes.

#### **CONCLUSION**

Threaded conversations are the core social data structure of the internet. While other, more technical data types, like DNS, have received significant investment, the ubiquitous “thread” has languished in contrast. Our effort is to invest in this data type by providing richer tools to visualize and manipulate them.

This paper offers two contributions for designers and practitioners of conversation-based environments, such as email, email lists, newsgroups, and instant messaging. First, we present design ideas for the creation of thread-based message browsers. While this paper describes our newsgroup visualization design, our ideas can be generalized to all message environments where the thread is the core data type. Second, we present a user study which demonstrates that thread-based browsers offer a superior user experience over message-based browsers, increasing user productivity, ease-of-use, and satisfaction primarily when dealing with complex threaded environments. These results suggest that future user interface designs of conversation-based environments should continue to focus

on the thread as the core data structure and work to support thread-centric tasks.

#### **ACKNOWLEDGMENTS**

We thank our user study participants, the contributors to the Usenet discussions we visualized, and the PSS and MSR CTG folks. Special thanks to Paul Johns, Duncan Davenport, Dany Rouhana, and Tony Capone.

#### **REFERENCES**

1. Boy, D., Lee, H., Ramage, D., Donath J. (2002) Developing Legible Visualizations for Online Social Spaces, *Proc. HICSS 2002*.
2. Donath, J., K. Karahalios and F. Viegas (1999) Visualizing conversations, *Proc. HICSS-32*.
3. Fiore, Andrew, Teirnan, S., Smith, M. (2002) Observed Behavior and Perceived Value of Authors in Usenet Newsgroups: Bridging the Gap, *Proc. CHI 2002*.
4. Lewis, D. and K. Knowles (1997) Threading electronic mail: A preliminary study, *Info. Processing and Management* 33(1), Elsevier, pp. 209-217.
5. Newman, P. (2002) Exploring discussion lists: Steps and directions. *Proc. JCDL 2002*, pp. 126-134.
6. Popolov, D., M. Callaghan and P. Luker (2000) Conversation space: Visualising multi-threaded conversation, *Proc. AVI 2000*, pp. 246-249.
7. Rao, R. and Card, S. K., Exploring large tables with the table lens, *Proc. CHI 1995*, pp. 403-404.
8. Rohall, S., D. Gruen, P. Moody and S. Kellerman (2001) Email visualizations to aid communications, *Proc. InfoVis 2001*, IEEE, pp. 12-15.
9. Sack, W. (2000) Conversation map: A content-based Usenet newsgroup browser, *Proc. IUI 2000*
10. Shneiderman, B. (1994) Dynamic Queries for Visual Information Seeking, *IEEE Software*, 11(6), pp. 70-77.
11. Smith, M., J. Cadiz and B. Burkhalter (2000) Conversation trees and threaded chats, *Proc. CSCW 2000*.
12. Smith, M. and A. Fiore (2001) Visualization components for persistent conversations, *Proc. CHI 2001*.
13. Venolia, G., and Neustaedter, C. (2003) Understanding Sequence and Reply Relationships within Email Conversations: A Mixed-Model Visualization, *Proc. CHI 2003*.
14. Viegas, F. and Donath, J. (1999) Chat circles, *Proc. CHI 1999*.
15. Xiong, R., Donath, J., (1999) PeopleGarden: Creating Data Portraits for Users, *Proc. UIST 1999*.