EVALUATING THE BENEFITS OF TILED DISPLAYS FOR NAVIGATING MAPS

Robert Ball, Michael Varghese, Andrew Sabri, E. Dana Cox, Chris Fierer, Matthew Peterson, Bill Carstensen* and Chris North
Department of Computer Science
Department of Geography*
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061, USA
email: {rgb6,north}@vt.edu

ABSTRACT
Maps are tools used by almost everyone in society for a variety of applications. However, when maps are used with computers they are almost always used with small, low pixel count displays, such as desktop monitors. We performed two experiments involving map usage with various tiled display configurations (one monitor, four monitors, and nine monitors). The first experiment focused on basic map navigation tasks and the second experiment focused on how to maximize the effectiveness of the details-on-demand interactive technique with large, high pixel count displays. We conclusively found from the experiments that finding objects and route tracing in maps was performed on average twice as fast on the nine monitors as the one monitor. We also found that participants on the nine monitor configuration had 70% less mouse clicks, 90% less window management, and a general accuracy and performance improvement over the one monitor. This indicates improved insight for large, high pixel count displays.

KEY WORDS
high resolution, maps, navigation, usability

1 Introduction

The motivation behind this experiment comes from evaluating the effectiveness of using large, high pixel count displays when navigating maps. Essentially, our research questions were:

- Does the use of large, high pixel count displays help navigate maps more efficiently?
- Can the higher pixel count that comes with tiled displays be used to improve performance with navigation techniques such as details-on-demand?

To answer our research questions we conducted two experiments involving large, high pixel count displays with maps. For each of the two experiments we used three monitor configurations: one, four, and nine tiled monitors. Figure 1 shows a person using the tiled display. The two experiments we used were: Navigating Large Maps and Using Details-on-Demand for Strategic Planning.

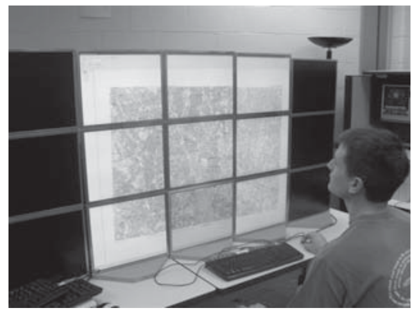

Figure 1. Example of navigating a large map on nine tiled monitors at a resolution of 3840x3072.

This paper is based on the results from Ball and North [3] which show that higher pixel counts positively affect basic navigation tasks with a static view. They show that people do not always perform better at higher pixels counts. Using one, four, and nine monitor configurations, they showed that if the target size is large enough people can zoom out to get an adequate overview of the visualization.

However, where their study had images that were the same size as their largest display, our first experiment used a map that was 50% larger than the largest display. The second experiment only had tasks directed towards the overview of a map, therefore, it used a "best-fit" overview for each monitor configuration without the need to zoom. Ball and North's results also were dependent on non-optimal image viewing software using basic pan and zoom techniques. As a result we used top of the line GIS software that has much better navigation techniques.

We wanted to know if other navigation techniques besides pan and zoom allowed for similar increases in performance. We chose a details-on-demand navigation technique as details-on-demand is fundamentally different from pan and zoom.

2 Previous Work

Different studies have been performed on large screens and multiple screens to compare their effectiveness to that of small or single screens. As mentioned above, the paper that most relates to this work is from Ball and North [3]. Other papers include generally usability of high pixel count displays from Ball [2] and Czerwinksi [5]. Tan et al. show how performance on a large screen can be better than a conventional screen even at the same resolution [13].

A range of research has been performed investigating usage with tiled displays. Two example field studies include [11] and [7]. Hutchings, et al. have performed a number window management experiments (e.g. citehutchings:revisiting:2004.

A few interaction techniques developed on multiple displays include pen-based approachs[8], mouse-based approaches [4], and head-tracking approaches [1].

3 Hardware Used

We assembled nine monitors that all run off one computer with 5 dual-head graphics cards. We created a 3x3 matrix of monitors with minimal effort. We also removed the plastic bezels that surrounds each monitor to reduce the distance between monitors.

Figure 2. The monitor configurations were one, four, or nine monitors with a resolution of 1280x1024, 2560x2048, and 3840x3072 respectively.

3.1 Protocol

All volunteers for the two experiments were screened prior to participation. All participants were required to have normal to corrected-normal vision, no color blindness, no familiarity with traveling/navigating through the state of Rhode Island (all maps used were from the state of Rhode Island), and no prior experience with large displays.

All participants were undergraduate students between the ages of 18 and 24. Twenty-four people participated in the navigating large maps experiment . They were all male computer science undergraduate students who received extra credit for their participation. Thirty-six participants were used for the details-on-demand strategic planning experiment. Twelve participants were female and twenty-four were male. Participants were 78% undergraduate computer science or computer engineering majors while the rest was a mix of different types of undergraduate majors. Sixty-six percent of the participants did the strategic planning experiment for extra credit.

All statistical analyzes for this paper were performed in SAS's JMP using standard ANOVA techniques.

4 Navigating Large Maps

Maps are used for a variety of reasons. More common usages include route tracing to more complex tasks such as deciding where building should be erected. As a result, our first experiment included six different types of tasks: Three search tasks; two route tracing tasks; two counting tasks; five comparison tasks (e.g. Which destination is closest); three intermediate tasks (e.g. Find the deepest water in Providence River); four advanced understanding task (e.g. Why is this area not developed?)

Using a between-subject design, all participants used the same map of Providence, Rhode Island. Performance time and accuracy were recorded as the dependent variables. Each participant was randomly assigned to a monitor configuration. Each participant was given a brief five to ten minute tutorial on how the software worked using a practice map prior to the actual experiment.

For this experiment, we used pan and zoom as our interactive technique. The software that participants used to navigate the map was ArcView: a full-featured GIS software program for visualizing geographical data by Environmental Systems Research Institute (ESRI). Participants were filtered to ensure they had no prior experience with ArcView.

4.1 Large Map Quantitative Results

We found the search task correlated to display size ($p < 0.01$) with differences between the one and four monitor configuration and between the one and nine monitor configurations. Looking at figure 3 one can see that the search task was performed more than twice as fast on the nine monitor configuration than on the one monitor configuration.

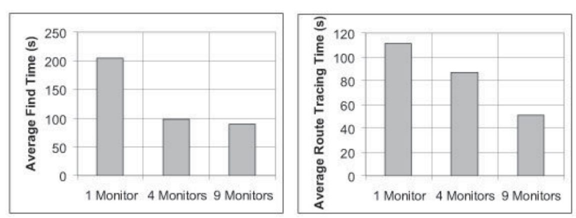

Figure 3. a) Average time in seconds for a participant to find a particular object or location on the map at different monitor configuration sizes. b) Average time in seconds for a participant to trace a route on the map at different monitor configuration sizes.

The route tracing tasks showed statistical significant ($p < 0.01$) when correlated to display size. Participants were instructed to trace a route from a source location to a destination location. Participants were able to accomplish this task on the nine monitor configuration more than twice as fast than on the one monitor configuration. Figure 3 shows the same trends as figure 3; participants were able to trace routes more than twice as fast on the nine monitor configuration compared to the one monitor configuration.

One reason for the increase in performance times could be the fact that participants on the larger monitor configurations navigated less with the map. In effect, as participants could see more of the map at a time, less navigation was required and consequently more time could be spent on the task at hand.

4.2 Large Map Qualitative Results

Although bezels are generally considered a distraction [10], we observed that participants used bezels to their advantage. A bezel is the border between monitors. A bezel is the limiting factor of how close two monitors can be together. When using the four and nine monitors participants would use the bezels to segregate the map into portions. By dividing the map into parts they were able to better keep track of which part of they map they had previously searched.

Similar to Furnas and Bederson's predictions [6], we observed that participants did not in general like to zoom in. If possible, participants would use the bounding box zoom rectangle to clip out all unnecessary parts of the map for the task. Then participants would often squint at the overview to try to gain as much detail as possible without having to actually zoom in any further and lose context of the entire overview.

We also observed that on the one and four monitor configurations people were more algorithmic in their approach to finding objects. As explained, on the nine monitor configuration participants rarely zoomed in. So, for most tasks, especially the search tasks, participants on the nine monitor configuration would use more intelligent heuristics to finding an object. For example, instead of searching the entire map for a university, as did participants on the one and four monitor configurations, the participants on the nine monitor configuration wold search logical areas, such as dense city areas or other areas that a university would logically be located. This might give indication of increased insight and awareness into the overall map [12]. By having a higher pixel count, participants were able to get a more accurate mental model of the map.

5 Details-on-Demand Strategic Planning

The main motivation for our details-on-demand experiment was to evaluate the results of [3] on a interactive technique that is distinctly different from pan and zoom for performance increases on large, high pixel count displays. Another motivation for this experiment was to see if the additional pixels could be used to increase the usability of the application by adding additional details to the overview.

There were two different versions of the experiment. The first version of the experiment did not display any details on the icon, just an image of soldiers. The second version included displaying a team icon with aggregated details about the team on the icon (see figure 4).

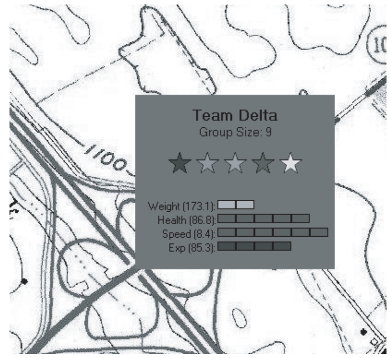

Figure 4. Example of an icon with details in the overview for the nine monitor configuration.

We presumed that the first version of the experiment had been created for the use of a single monitor. The first version did not have any details on the icon because it is hard to have an icon that is both small enough for the underlying map to be usable and have readable details on the icon itself. The second version was created with the intent of changing the first version to take advantage of the extra pixels that come with tiled displays.

The amount of detail on the icon displayed in the second version was dependent on the monitor configuration size. While maintaining the same area of the map, icons that had more pixels had more details on them. This idea is shown in figure 5. Figure 5 shows how the icons in the left image take up as much area as on the right image even though the left image is a screen shot of the nine monitor configuration and the right image is a screen shot of the one monitor configuration.

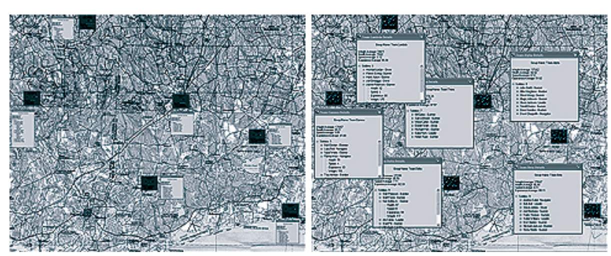

Figure 5. a) Screenshot from the nine monitor configuration. b) Screenshot from the one monitor configuration. Both screenshots have been shrunk to show that the icons take up the same area in the map. The screenshots are not proportional to each other.

The reasoning behind this design decision was based on the fact that more pixels were available for each icon at larger monitor configurations. By keeping the size of the icon to the size of the map ratio the same (the same area), the larger the monitor configuration, the more room was available for aggregated details in the overview.

5.1 Experiment Setup

The experiment consisted of displaying a map with multiple army teams located across the map. The team locations were represented by a team icon, with the lower left hand corner representing the team's exact location. A left click on a team icon displayed a popup window about the team and their statistics (see figure 6). A range of team information was displayed within the window (e.g. team name and average health). Listed with each soldier's name was their respective job and personal statistics. The window was resizable and movable.

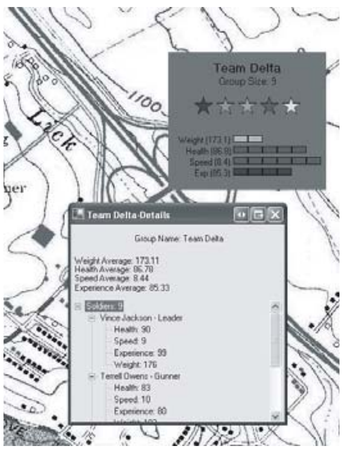

Figure 6. Example of an icon with detail in the overview with an associated popup window showing details.

The experiment was developed using C# in a Visual Studio.NET environment with Microsoft Access as a backend. All participant moves, resizes, and clicks were tracked and recorded by the application. In order to change scenarios and enable or disable monitors quickly UltraMon was used.

Participants were told that they were to act as a general in an army to decide which team should be used for an ensuing battle. Participants had a range of tasks which included different types of analysis. The following list shows the different types of tasks: First, details only: Find the weakest soldier. This task entailed looking at the details of every soldier. The aggregated details did not help in this task. Second, overall statistics of teams: Find the team with the strongest overall statistics. Third, map only: Find the team with the best geographic location in reference to the attack point. Fourth, map and overall statistics of teams: Find which team is the *best* to attack with based on the best geographic location and statistics. Fifth, map and overall statistics of teams: Find which team is the *worst* to attack with based on the worst geographic location and statistics.

In this experiment we used three different monitor configurations as the first experiment (one, four, and nine tiled monitors). However, we used a within-subject design so that every participant used every monitor configuration. However, participants only used one version of the experiment (with or without aggregated details) for a between-subject design. We used two Latin Square designs (one for each of the two versions) for counter balancing. To get an equal distribution of gender, we used six females and twelve males for both versions of the experiment.

We also used three different scenarios each with a different map. We did this so that participants would not become overly familiar with a single map. All scenarios were performed in the same order. A ten minute tutorial was given to the participant about how to use the application in order to baseline the participants.

In summary, our independent variables were: Monitor configurations (one, four, or nine monitors); with or without aggregated details in the overview; map (three different maps from Rhode Island). Our dependent variables were: amount of interaction (clicks, moves, and resizes), performance time, and accuracy.

5.2 Interaction Quantitative Results

By recording mouse events of each participant, we were able to track how much interaction took place. We found that the monitor configuration size statistically correlated to the number of mouse clicks ($p < 0.001$). We also found that our independent variable of with or without aggregated details also correlated to mouse clicks ($p = 0.003$). We tested for an interaction effect and for multicollinearity and did not find either.

Looking at figure 7 one can see that participants on the nine monitor configuration clicked on average 70% less than on the one monitor configuration. Also, participants that had aggregated details in the overview clicked 15% less than participants that did not have the extra details. All that can be concluded about the mouse clicks is that more mechanical effort was required for the smaller displays.

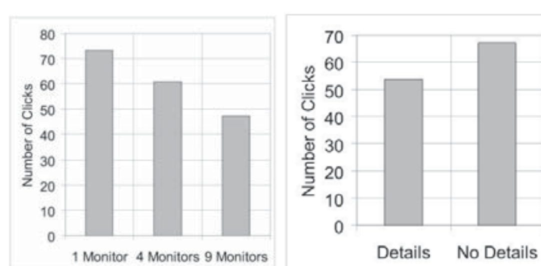

Figure 7. a) Chart showing the trend of number of clicks decreasing as monitor configuration increases. b) Chart showing how the number of clicks was less with icons that had aggregated details.

We found similar results with window moves (i.e. window management). Participants often moved the popup window for a number of reasons. For instance, they moved windows closer to compare different team statistics and they moved windows to see the underlying map. We found that the monitor configuration size statistically correlated to the number of mouse moves (p = 0.0509) as did the aggregated details variable (p = 0.0572). Again we tested for an interaction effect and for multicollinearity and did not find either.

The number of window moves was approximately equal for the four and nine monitor configurations. However, the four and nine monitor configurations moved windows approximately 90% less than the one monitor configuration. We also found that participants that used the aggregated details in the overview moved windows 60% less than than participants that did not used the aggregated details in the overview.

This data objectively supports many of the subjective finding in [2] and [9] that shows that people perform less window management with higher pixel counts. Also, the implications of this decrease in window management means that people can use tiled displays to help them perform less window management and focus more on the tasks they want to accomplish.

5.3 Performance and Accuracy Quantitative Results

The users on the nine monitor configuration performed better on the second task, which asked about overall statistics, and had an interactive effect between monitor configuration size and the details variable (p=0.04). The fourth task, finding the team with the best overall statistics and geospatial location had similar results as the second task with an interactive effect (p = 0.002).

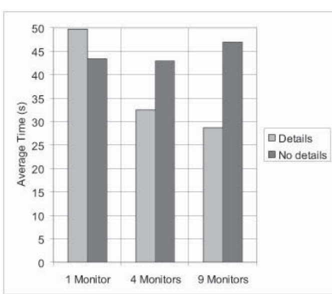

Figure 8. Chart showing general trends for finding the team with the best statistics.

Figure 8 shows the average time to complete the second task did not vary much for the non-details version of the experiment as the monitor configuration size increased. However, for the details version, as more details were added as each monitor configuration size afforded more pixels, the task performance was improved. For the aggregated details version of the experiment, participants performed 57% faster on the nine monitor configuration compared to the one monitor configuration.

Accuracy was found to be statistically correlated to the aggregated details variable (p = 0.0115), but not to the monitor configuration size. Participants who used the details version of the experiment totaled 47 incorrect answers, while participants who used the non-details version, totaled 74 incorrect answers.

This trend of greater performance, accuracy and the fact more details can be shown on larger configurations indicates that these large configurations can improve insight and understanding the task at hand.

5.4 Qualitative Results

Window management was a major issue between the different monitor configurations that participants were confronted with. When participants were given more monitors to work with, they did not necessarily use it to their advantage. Many participants during the study would move their windows to different locations. One common technique participants used on larger monitor configurations is putting all team windows together on the same monitor. This method of clustering windows together was consistently seen as participants were given more monitors.

Figure 5 shows much of the screen the popup windows used on the different monitor configurations. The left image shows the amount of room that the windows take up on the nine monitor configuration while the image on the right shows the amount or room that the windows take up on the one monitor configuration.

A map reading error that participants often made was focusing on a specific area of the map and not observing other areas. This was especially prevalent on the nine-monitor configuration where participants would focus their attention on where most teams were and not analyze teams that were farther away.

Participant's responses indicate that participant preference for monitor configuration size was split. We found that 15 of the 36 participants preferred the four monitor configuration. The four-monitor configuration provides somewhat of a threshold where participants can perform better than the one-monitor configuration without the overwhelming setup of nine monitors. Sixteen participants preferred the nine-monitor configuration. Participants cited the novelty of the nine-monitor configuration as the reason for their preference. Three participants preferred the one-monitor configuration. Their reasoning was that the larger-configurations were overwhelming. The remaining two participants had no preference.

6 Conclusion

We found a number of benefits for the nine monitor configuration compared to the one monitor configuration which include: Finding objects twice as fast; performing route tracing twice as fast; 70% less mouse clicks; 90% less windows management. With our details-on-demand strategic planning experiment we found that having more pixels allowed us to put aggregated details in the overview even with keeping the same ratio of icon size to map size. By putting aggregated details into the overview participants not only performed faster but 73% more accurately.

Overall, improved performance on geospatial interpretive tasks indicates that the large, high pixel count displays can improve insight and understanding into the data. With a larger screen real estate, people are able to analyze more data, more accurately, faster, and with less interaction. The results from this paper show that larger configurations can be applied to a variety of data analysis tasks.

Not only can the information visualization community benefit by being able to display more data and consequently interpret the data more accurately, these displays can also be used for a diverse set of applications including company analysis, government intelligence agencies, and even home computing.

7 Future Work

This paper has brought up a number of interesting issues that need to be addressed. The following is a list of future work:

- Improving interactive techniques and notification systems for large, high-resolution displays.
- More in-depth studies with maps and large displays using expert geographers and cartographers.
- Objectively measure the awareness and interaction advantages of large, tiled displays.

8 Acknowledgements

This research is partially supported by the National Science Foundation grant #CNS- 04-23611. This study was also supported and monitored by the Advanced Research and Development Activity (ARDA) and the Department of Defense. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the National Geospatial-Intelligence Agency or the U.S. Government.

References

[1] Mark Ashdown, Kenji Oka, and Yoichi Sato. Combining head tracking and mouse input for a gui on multiple monitors. In *Extended abstracts of CHI '05*, pages 1180–1191, 2005.

[2] Robert Ball and Chris North. An analysis of user behavior on high-resolution tiled displays. In *Interact 2005 Tenth IFIP TC13 International Conference on Human-Computer Interaction*, 2005.

[3] Robert Ball and Chris North. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *Extended abstracts of CHI '05*, pages 1196–1199, 2005.

[4] Patrick Baudisch, Edward Cutrell, and George Robertson. High-density cursor: A visualization technique that helps users keep track of fast-moving mouse cursors. In *Proceedings of Interact 2003*, pages 236 – 243, Zurich, Switzerland, 2003.

[5] Mary Czerwinski, Greg Smith, Tim Regan, Brian Meyers, George Robertson, and Gary Starkweather. Toward characterizing the productivity benefits of very large displays. In *Proceedings of Interact 2003*, 2003.

[6] George W. Furnas and Benjamin B. Bederson. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of CHI '95 Human Factors in Computing Systems*, pages 234 – 241, 1995.

[7] Jonathan Grudin. Partitioning digital worlds: Focal and peripheral awareness in multiple monitor use. In *Proceedings of CHI 2001*, pages 458 – 465, 2001.

[8] Francois Guimbretire, Maureen Stone, and Terry Winograd. Fluid interaction with high-resolution wall-size displays. In *Proceedings of UIST 2001*, pages 21–30. ACM, 2001.

[9] Dugald Ralph Hutchings and John Stasko. Revisiting display space management: Understanding current practice to inform next-generation design. In *Graphics Interface 2004*, pages 127 – 134. Canadian Human-Computer Communications Society, 2004.

[10] Jock D. Mackinlay and J. Heer. Wideband displays: Mitigating multiple monitor seams. In *Proceedings of CHI '04*, pages 1521 – 1524, Vienna, Austria, 2004.

[11] Meredith Ringel. When one isn't enough: An analysis of virtual desktop usage strategies and their implications for design. In *Extended Abstracts of CHI 2003*, pages 762 – 763. ACM Press, 2003.

[12] Purvi Saraiya, Chris North, and Karen Duca. An insight based methodology for evaluating bionformatics visualization. *IEEE Transactions on Visualizations and Computer Graphics*, 11(4), July/August 2005.

[13] Desney Tan, Darren Gergle, Peter G. Scupelli, and Randy Pausch. With similar visual angles, larger display improve spatial performance. In *Proceedings of CHI '03*, April 2003.