Weave: A Web-based Architecture Supporting Asynchronous and Real-time Collaboration

Andrew Dufilie
UMass Lowell
1 University Ave
Lowell, MA 01854
adufilie@cs.uml.edu

John Fallon
UMass Lowell
1 University Ave
Lowell, MA 01854
John_Fallon@student.uml.edu

Patrick Stickney
UMass Lowell
1 University Ave
Lowell, MA 01854
patrick_stickney@student.uml.edu

Georges Grinstein
UMass Lowell
1 University Ave
Lowell, MA 01854
grinstein@cs.uml.edu

ABSTRACT
This paper presents Weave, an open-source application framework designed to support both asynchronous and real-time collaborative visualizations. While other frameworks add support for session state as an afterthought, Weave was designed from the beginning with sessioning in mind. The framework allows run-time linking and modification of any parameter. Changes in the session state are stored in memory to allow undo, redo and replay capabilities. A session history file can be saved and shared with other analysts for asynchronous collaboration. Weave can also share live session-state updates over a chat server to allow real-time collaboration across multiple web browsers, whether the analysts are co-located or not. Sharing not just the visualization and data, but also the process involved to arrive at the visualization provides numerous opportunities for further research in collaborative visual analytics for the masses.

Categories and Subject Descriptors
D.2.11 [**Software**]: Software Architectures – *Domain-specific architectures.*

General Terms
Design.

Keywords
Visualization, Session History, Collaboration.

AVI '12, May 21-25, 2012, Capri Island, Italy

1. INTRODUCTION
Recently, web-based data visualizations have become commonplace, providing accessible visual representations of data for various purposes to public audiences with diverse backgrounds. Many of the visualizations found on the web are one-off implementations customized for only one set of data. Other systems provide users with the ability to upload and explore their own data and asynchronously collaborate through annotated visualizations. However, none of these systems provide real-time collaborative visualization and exploration of arbitrary data sets across the web. In this paper we present the data visualization platform Weave, which addresses these needs. The general goal of the Weave project is to create a scalable web-based platform that supports data visualization, exploration, analysis, session history and real-time collaboration. Baumann et al.[1] provide an in-depth discussion about the design goals of Weave related to session history and collaboration.

2. RELATED WORK
There is a large amount of related work in collaboration with respect to chat rooms, virtual worlds and collaborative workspaces (see the CSCW Conferences [2]) but none with respect to shared interactive visualization systems with history support for both private and shared spaces. The classic collaboration is via shared screens within which the cursor (or control) is passed (such as Webex or VPNs). Some systems provide support for asynchronous collaboration with both data and visualizations, but the visualizations are the end result of the analysis and are typically not interactive nor do they contain any history.

Many Eyes, from IBM, provides support for one form of asynchronous public collaboration.[4][7] In the Many Eyes environment, all data used during collaboration is made public and any visualizations created are posted on the internet for public comment. Although useful for exchange of ideas, this form of collaboration offers little to no data privacy and does not enable real-time interactions.

General Dynamics C4 Systems developed Command Post of the Future for the U.S. Department of Defense [3]. This collaborative application, which supports synchronous and asynchronous collaboration, is used in the United States military. Its main purpose is to visualize a battlefield and let soldiers of different ranks communicate data and tactics. This is not an open-source project and it is only for military use.

VizCept [8], a web based visualization system which was presented at IEEE VisWeek 2010, has support for real-time collaboration. This system is designed entirely for text-based visualization and analysis. While it has workspaces for private use and shared use, the data being used is available to all participants of the collaboration session.

3. WEAVE

The evolution of Weave has been supported and guided by members of the Open Indicators Consortium (OIC) [6]. The common desire of all consortium members is to build an open-source visualization framework to allow anyone to share and analyze data, but the plans of how to use that framework differ greatly. The flexibility of Weave has evolved as a result of the cooperation of the consortium, and without the guidance of this diverse group Weave would not be what it is today.

3.1 Session-driven Framework

The Weave Framework was designed to support highly interactive, linked visualizations for users of varying skill levels. To support the desired level of functionality and customizability, all visualization components in Weave are built on top of a core sessioning framework which allows linking variables to the user interface. Any variable appearing in the session state can be changed at run-time to reflect a change in the interface and similarly any change in an interface variable is reflected in that session state. The framework also supports end-user programming for a highly customizable interface.

3.2 Asynchronous Collaboration

Weave supports asynchronous collaboration. Weave does this through its innovative session state and session history functionality. The session history within Weave keeps track of all changes made during the creation of a visualization. This session history is exportable to a .weave archive that holds the entire session for Weave. This file can be shared with anyone, and then loaded into a Weave instance. Once loaded into Weave, the previous state will be restored and the entire history leading up to the current state will be accessible by the new user. The history can be played back to see how the current state was reached.

Session history allows study of the thought process that leads up to the creation of a visualization. The user may add annotations for each step taken in creating a visualization and, when replayed, these annotations can appear within each step of the session history. As mentioned earlier, Many-Eyes is a system that supports asynchronous collaboration. With Many-Eyes a visualization made by someone can be viewed and then modified later, but there is no access to the process of how the creator arrived at the resulting visualization. Weave allows exploration of the visualization creation process. This can for example be used to provide training support and help novice users become more proficient at generating visualizations within Weave. An example of Weave's first released session saving ability at work can be found at the Metro Boston Data Common website [5].

3.3 Synchronous Collaboration

Weave's collaboration support does not stop at asynchronous collaboration, it also supports synchronous collaboration. Unlike the asynchronous aspect of Weave, real-time collaboration can take place involving multiple people editing and working on one synchronized state that all the participants share. The participants can be co-located and can be using different devices, or they can all be remotely participating in a collaboration session. This is achieved through the use of a Jabber server to send and receive real-time messages. See Figure 1 for an example collaboration session.

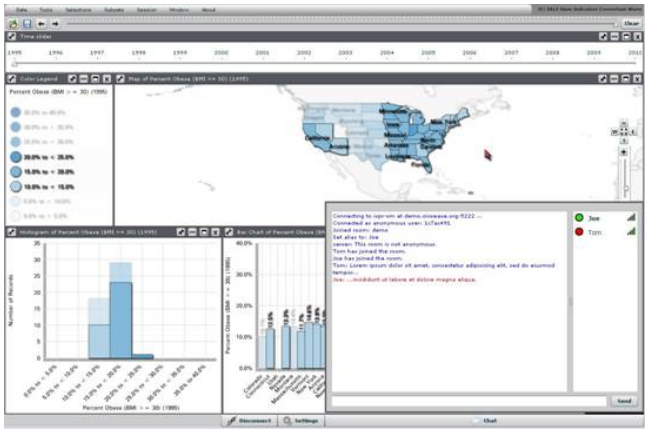

Figure 1. Weave collaboration interface.

To maintain efficiency whenever a change in the session state is made, the forward and backward differences are calculated and sent to the chat room rather than the entire state. The session diff is used instead of the entire session state so that the changes can be made on everyone's systems with as little network traffic as necessary. For quick parsing and transmission, Weave uses the binary AMF3-encoding of objects rather than using text serialization. Since the Jabber protocol only supports text, the objects must be encoded in base64 before transmission.

One of the issues with collaborating in real time is lag. Lag is a drawback for all applications that rely on a network to communicate with other people. If every interaction required permission from a central server, the interface would seem sluggish. Weave uses session diffs to overcome this problem. The session diffs contain identifiers indicating the order they occurred locally. If two users make a change at the same time, the server will broadcast them in a particular order. The user whose diff appears later will roll back and re-apply the diffs in the order the server has broadcasted them. This method allows each user to have a responsive interface during real-time collaboration.

Voice and video chat are supported within Weave. The RTMP protocol was used to implement both the voice and video streaming service. The RTMP protocol was a good choice since Weave is a flash-based application, but other protocols designed for streaming could be used to support real-time collaboration. An open-source server designed to manage the RTMP protocol is used to support interactions between each instance of Weave that is collaborating.

4. FUTURE WORK

While collaboration in Weave is functional, there are still multiple areas in the interface which can be improved. The user interface used for collaboration right now is a prototype and rigid. In the future, it will be made more flexible so that the chat interface can be separated from the visualization interface to provide more screen space for building visualizations.

Weave also opens up a large potential for future research related to session history and collaboration. When navigating the detailed session history, one may find it hard to pick out the important parts of the analysis process. We will investigate the possibilities of having the system automatically determine the major steps in the analysis process and we will provide a way for the user to manually flag bookmarks as significant points in the session history. Possible applications for automatic detection of these key states would be to compare thumbnails of each session state to determine large visual differences, or to deem certain actions as major actions, such as opening or closing a visualization panel. The interface for navigating the session history could be made similar to the thumbnail popup when the mouse is hovered over the timeline in a YouTube video, or a series of thumbnails as done in PowerPoint.

In addition to allowing a way of navigating to key states or bookmarks in the session history, a dynamic query interface is planned to allow the user to cut down the history into fewer steps based on tags associated with the objects referenced by the diffs. For example, if the user is only interested in changes related to a particular bar chart, they could enter a query whose results would show a list of the steps in the session history related to that bar chart. A user could also search the session history for a particular data attribute to see where in the analysis that attribute was investigated.

5. CONCLUSION

We have developed a foundation for a general web-based application framework with broad, expanding goals in mind to enable data visualization, exploration, analysis, session history and real-time collaboration. Weave's focus on sessioning from the beginning of its development has aided greatly in creating a framework that supports collaboration. The collaboration aspect of Weave works asynchronously and synchronously to tailor to each user's specific needs. Weave is a large-scale open source platform and web-based visualization system that supports collaboration between multiple people that can be co-located or remotely located. The project opens up many possibilities for analysts and researchers alike.

6. ACKNOWLEDGMENTS

The Weave project is funded by members of the OIC, the Knight Foundation [9] and the University of Massachusetts Lowell. Weave is supported by a twenty-member software and data development team led by University of Massachusetts faculty members, Dr. Georges Grinstein and Dr. William Mass and involving numerous Computer and Social Science masters and doctoral students. For more information, visit www.oicweave.org.

7. REFERENCES

[1] Baumann, A., Grinstein, G.G., and Mass, W.: Collaborative Visual Analytics with Session Histories. Unpublished Research Paper. No. 2009-005, University of Massachusetts Lowell, Dept. of Computer Science, Lowell, MA 01854.

http://teaching.cs.uml.edu/~heines/techrpts/details.jsp?Year=2009&SeqNo=005.

[2] CSCW Conferences. http://www.cscw2011.org/.

[3] General Dynamics C4 Systems. Command Post of the Future. http://www.gdc4s.com/content/detail.cfm?item=2a58f8e2-ef2b-4bb1-9251-42ee4961dd7f.

[4] Many Eyes. http://www-958.ibm.com/software/data/cognos/manyeyes/.

[5] Metro Boston Data Common. http://metrobostondatacommon.org.

[6] Open Indicators Consortium. http://bugs.oicweave.org/projects/weave/wiki/About_The_Open_Indicators_Consortium.

[7] Viégas, Wattenberg, van Ham, Kriss and McKeon (2007), "Many Eyes: A Site for Visualization at Internet Scale", IEEE Infovis Conference, Computer Society Press.

[8] VizCept. http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5652932.

[9] Fest, M.: Knight Foundation Spurs New Round of Local News and Information Projects Nationwide. Press release. Knight Foundation. January 13, 2010. http://www.knightfoundation.org/news/press_room/knight_press_releases/detail.dot?id=355583.