

Collaborative Multitouch Log Browsing

[Extended Abstract]

Jeff Wilson and Robert Biddle

School of Computer Science
Carleton University, Ottawa, Canada

jppwilso@connect.carleton.ca, robert_biddle@carleton.ca

ABSTRACT

System logs contain much information that can assist administrators in monitoring the state and history of system services. Tools and alarms that analyze such logs, however, are designed to emphasize certain patterns, and therefore make it difficult to detect novel problems or attacks. In this paper we describe a multitouch visualization environment to facilitate situational awareness of log data. By supporting administrators in identifying emergent patterns, we leverage the same skills they currently apply to crude text commands. Moreover large multitouch displays allow the environment to act as an information radiator and support collaborative exploration.

Categories and Subject Descriptors

H5.2 [Information interfaces and presentation]: User Interfaces—Graphical user interfaces; K.6.5 [Management of Computing and Information Systems]: Security and Protection

Keywords

Log visualization, Multitouch, System administration

1. INTRODUCTION

The apparent simplicity of web server logs conceals a challenging complexity of events. While many system administrators are experts at drilling into logs with simple textual tools and with languages like *awk* and *sed*, new technologies may support interaction that affords new benefits.

With the coming availability of large and affordable multitouch displays we are offered an opportunity to revisit visual languages in the exploration of data. We expect the use of touch will foster an *actual* rather than simply metaphorical sense of data manipulation and we hope that this immediacy will invite greater exploration. This capability will also improve communication of discoveries to a broad range of stakeholders.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VIZSEC '10 Ottawa, Ontario, Canada

Copyright 2010 ACM 978-1-4503-0013-1/10/09 ...\$10.00.

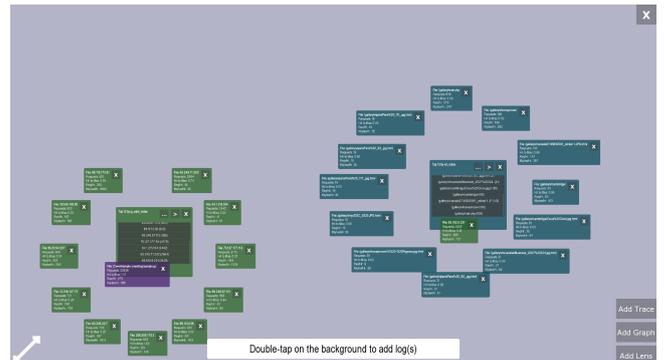


Figure 1: Expanding by IP then by URL.

Visualization has been used for system management for some time. Web traffic performance and e-commerce site usage have been common subjects (e.g. WebViz[7]), and other systems have supported security tasks such as intrusion detection [5]. Our emphasis is on situational awareness rather than supporting specific intentional tasks. In particular, we are motivated by Cockburn's idea for team-room *information radiators* [3], together with the direct manipulation potential demonstrated in *starfields* [1], and leveraging co-located collaborative inspection [8]. Our goal is to develop a platform for log visualization research using this approach.

Cockburn describes an information radiator as "*information in a place where passersby can see it.*" [3] (p.114) and this can take many forms, such as a frequently updated poster in a high traffic hallway. What is of particular interest is the change of status. When considering security questions, allowing idle manipulation of an ever-present status display may set in motion creative forms of speculation. Cockburn's approach is used in Agile software development, where this collaborative speculation leads to novel solutions well aligned to the needs of the application domain.

2. DESCRIPTION

We are developing a prototype log browsing tool (see screenshot Figure 1) using PyMT[4]. The design assumes the availability of a large work surface. We are using a multitouch display assembled from low-cost components using a diffuse illumination design, but our architecture and implementation also work with the flat wall-mounted displays that are becoming available.

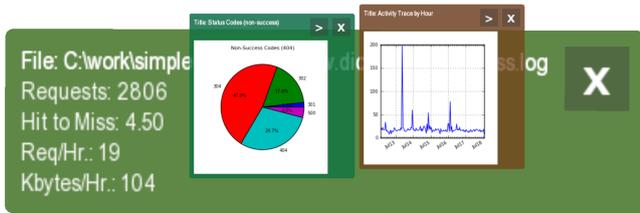


Figure 2: Graphical view.

PyMT supports input from a number of multitouch sources including TUIO, trackpad, and touchscreen HID devices. That made it a perfect fit for our target platform and also allowed for development and testing using other devices. Being Python-based, PyMT yielded the expected benefits of platform independence, rapid development turnaround, and a rich selection of extension packages. For a graphing extension we chose matplotlib. For statistics we are exploring scipy and rPy.

Users may open historical logs from the file system and these logs appear as semi-opaque rectangular objects showing basic descriptions in text. The system is also capable of processing and displaying live streams. Each object on the desktop is easily moved, rotated, and resized using simple touch gestures and the workspace supports simultaneous scaling of all objects. Inspection of objects is performed with configurable lenses[2]. With very few such tools the user is capable of navigating, filtering, aggregating, and graphing multiple data sources.

As an example of a session, in fig. 1 the user has opened a single log, shown in purple. Fig. 2 shows what happens to a visually scaled log when you hover over it with one or more graphing tools. Back in fig. 1, we can assume the user was curious about IP addresses. They hovered over the log with a lens (the square object inside the ring) configured for IP and clicked *expand* to create the green cluster of top-12 requesters. They then moved a subset object to the right using a gesture much like you see in the live session in fig. 3 and expanded it using a lens configured for URL. The result is that the rightmost grouping shows the top-12 resources requested by the top IP address from the first “query”.

If the graph tool (fig. 2) were superimposed over a stack of log objects the calculation would be present the union of the underlying datasets.

3. DISCUSSION AND CONCLUSIONS

We find that with few tools there is still remarkable flexibility in this system. Large sets of data can be sliced into smaller ones and then these in turn can be easily joined with other subsets for graphing. Without use of any kind of query language, we quickly found the most active resources on a web site for the previous week, filtered out the search-bot traffic (the IP addresses were consistent with the agent), and then saw that the most active requester was posting comments to a particular gallery. We could also see that the search bots were a reasonable burden on the limited bandwidth of the hobby site, that one specific PHP application generated the majority of site errors, and that there were some surprising geographical centres of interest.

While none of this is beyond a system administrator or DBA the interesting feature was that no scripting or cod-

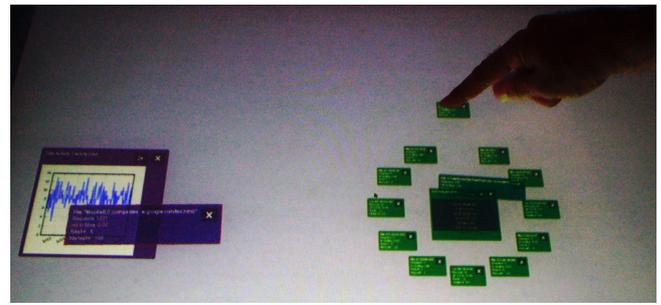


Figure 3: Live session.

ing was required. Moreover, we have found that the large display and the multitouch capability compels and rewards exploration. Even more encouraging is that the technology appears to encourage and support collaborative inspection in the way we had hoped.

In terms of future work, we are interested in the potential for reuse and sharing of discovered data patterns and so we will seek to codify the query implied by a user’s manipulations, creating what Nardi[6] refers to as a visual formalism. We are also interested in giving some lenses more sophisticated statistical and data mining capabilities with the use of plug-ins. We will be investigating a default view for real-time data for use as a proper information radiator. Finally, we will extend support from web logs to logs of other system services.

4. ACKNOWLEDGMENTS

We thank NSERC for an Undergraduate Student Research Award, NSERC ISSNet and SurfNet for research funding, and Don Wilson & Associates for sample analysis data.

5. REFERENCES

- [1] C. Ahlberg and B. Shneiderman. Visual information seeking: tight coupling of dynamic query filters with starfield displays. In *CHI '94*, New York, 1994. ACM.
- [2] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: the see-through interface. In *SIGGRAPH '93*, New York, 1993. ACM.
- [3] A. Cockburn. *Agile Software Development: The Cooperative Game (2nd Ed.)*. Addison-Wesley, 2006.
- [4] T. E. Hansen, J. P. Hourcade, M. Virbel, S. Patali, and T. Serra. Pymt: a post-wimp multi-touch user interface toolkit. In *ITS '09*, New York, 2009. ACM.
- [5] A. Komlodi, J. R. Goodall, and W. G. Lutters. An information visualization framework for intrusion detection. In *CHI '04*, New York, 2004. ACM.
- [6] B. A. Nardi. *A Small Matter of Programming: Perspectives on End User Computing*. MIT Press, Cambridge, MA, USA, 1993.
- [7] J. Pitkow and K. A. Bharat. Webviz: A tool for world-wide web access log analysis. In *WWW '94*, 1994.
- [8] S. D. Scott, K. D. Grant, and R. L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW'03: Proceedings of the eighth conference on European Conference on Computer Supported Cooperative Work*, pages 159–178, Norwell, MA, USA, 2003. Kluwer Academic Publishers.