

;login:

THE MAGAZINE OF USENIX & SAGE

June 2001 • Volume 26 • Number 3

inside:

USENIX GOOD WORKS

Research Exchange Program Update from
the Field: Vrije Universiteit and the
University of California, San Diego,
Cooperative Association for Internet
Data Analysis

by Wilfred Dittmer

USENIX & SAGE

The Advanced Computing Systems Association &
The System Administrators Guild

Research Exchange Program Update from the Field

by Wilfred Dittmer

Scientific Programmer, Vrije
Universiteit, Amsterdam,
Netherlands
<wtrittmer@cs.vu.nl>

A report on the ReX exchange program jointly supported by USENIX and NLnet. See <<http://www.usenix.org/about/rex.html>> and <<http://www.nlnet.nl/projects/rex/>> for information about this valuable program.

Vrije Universiteit and the University of California, San Diego, Cooperative Association for Internet Data Analysis

USENIX and NLnet offered me the possibility to work for the Cooperative Association for Internet Data Analysis (CAIDA) in San Diego, CA, for a period of six months via the Research Exchange (ReX) program. At the time, I had just graduated and was going to work for the Vrije Universiteit (VU) in Amsterdam. VU didn't have a job opening, so it was great to participate in the exchange in the meantime. The purpose of the exchange was to enhance and develop tools for measuring Internet traffic in order to obtain better insight into the physical organization of the network through the Skitter project at CAIDA. First, I will tell a bit more about VU and CAIDA.

VU <<http://www.vu.nl>> is where I graduated with a degree in computer science. As a student, I worked on the GLOBE <<http://www.cs.vu.nl/~steen/globe>> project, which is a novel scalable infrastructure for a massive worldwide distributed system. I wrote an application that uses GLOBE to locate and contact users over the Internet, called Loc8 <<http://www.cs.vu.nl/~baggio/loc8.html>>.

CAIDA <<http://www.caida.org>> is an organization that uses its skitter tool to actively probe the Internet in order to analyze topology and performance. Skitter measures forward IP paths by recording each hop to many destinations. Using ICMP echo requests, skitter also collects Round-Trip Time (RTT) to those destinations. The data collected provides indications of low-frequency, persistent routing changes, and correlations between RTT and time of day may reveal a change in either forward- or reverse-path routing. The skitter data can also be used by their Otter tool to visualize the directed graph from a source to much of the Internet. Otter can handle visualization tasks for a wide variety of Internet data, including datasets on topology, workload, performance, and routing. CAIDA has written many more tools to collect, analyze, and visualize data: for example, CoralReef and Walrus. Now on to the projects I worked on for CAIDA.

The first project I worked on was to map reverse-traceroute and looking-glass servers onto a world map. Using Internet search engines and emails sent by people maintaining reverse-traceroute and looking-glass servers, I collected as many reverse-traceroute and looking-glass pages as I could. Using IP addresses and NetGeo, I could find their latitude and longitude as well as city, state, and country. With this information I then used Geo-Plot to display the servers on a world map.

The servers appear as little dots on the map; clicking on them brings up a page containing the URLs for the traceroute and looking-glass servers for that node or will zoom in to that particular region. The nodes are colored to indicate what services are available for that location.

The information collected was also used to obtain the AS numbers for all the traceroute and looking-glass pages (again using NetGeo) and to build a Web page that allows users to search our database for traceroute and/or looking-glass servers by AS number. This page also allows users to search by AS name, latitude, longitude, city, state, or country.

The purpose of these pages is to give users with network problems the ability to find the nearest traceroute/looking-glass service, in a convenient way based on the location they desire to trace from. The traceroute/looking-glass services can then be used to view their own network and can provide hints to what causes their networking problem. The page can be found at: <http://www.caida.org/analysis/routing/reversetrace>.

The second project was to enhance the Web pages that provide users with daily summaries of data collected by the skitter boxes around the world. I first had to familiarize myself with the way the scripts currently processed the data for Web usage before I was able to expand their Web pages. In particular, I added pie charts to display the percentage of paths going through a particular country or AS. Because a user can select multiple skitter boxes and dates to generate this pie chart, I had to merge the data of the skitter boxes and dates before I could generate the pie charts.

I also added AS connectivity graph images to the Web pages. An AS graph image shows AS nodes arranged in a circle with polar coordinates, with their angle based on the location of their AS headquarters and their distance from the center reflecting the richness of their observed connectivity to other ASes. Those nodes closest to the center are the most connected ASes as observed for those particular days and skitter boxes. The graph is drawn by Otter. Users can download a data file for Otter that allows the user to view the graph on their own machine with more detail. The page can be found at: http://www.caida.org/cgi-bin/skitter_summary/main.pl.

The third project was to enhance the Otter network mapping/visualization tool so that it has arrows at the end of the links between nodes to indicate the direction of the link. We would want to print the arrows too, which required a better understanding of how PostScript works, since Otter saves images in PostScript format. After this, I implemented the arrangement of the links entering or leaving the node. Instead of all links merging in the center of the node, the lines are now spread out over the four sides (a node is displayed as a square) depending on the location of the node it is connected to. Links are also sorted by up/down or left/right to the same node.

The final project, which I was unable to finish because of lack of time, was generating animations of AS graph images for specific servers aggregated over a particular number of days. I used Otter to generate the graphs. Otter will use the maximum size of a page when it's printing to file. For animations this is not preferable, because the images will vary in size and common nodes between frames will jump over the screen. So we define four anchor points in every frame, with the minimum value set to the minimum outdegree found in all frames. After this we normalize all outdegrees to be a percentage of the maximum outdegree found in all frames.

I extended the program to use only a limited number of nodes in the animation to improve visibility of the most connected nodes. After doing a day-by-day animation for 364 days in 2000, the nodes were still jumping a little. The way to solve this problem is to generate intermediate frames that smooth out the movement of the nodes between frames. Because of time constraints, this has not yet been implemented.

Future projects at CAIDA include the integration of bandwidth estimation data into a database with Otter and/or Walrus as a front end. The data will be collected by a tool suite called netchar that uses features from Bruce Mah's pchar, Allen Downey's clink, and Constantinos Dovrolis' pathrate. The data-processing back end refines and extends capabilities in Allen Downey's clink, estimates the bandwidth and latency of each link appearing in a path, and will assemble a view of the subnets connecting the hosts.

REFERENCES TO PROGRAMS

USED/MENTIONED:

Reverse traceroute servers allow the user to traceroute from the server's perspective. The user can see the return path from that server to his or her own machine or to other machines.

Looking-glass servers allow users to see BGP tables at a particular point in the Internet.

NetGeo uses "whois" records to determine the location given a domain name. See <http://www.caida.org/tools/utilities/netgeo>.

Geoplot is a light-weight java applet that allows users to create a geographical image of a data set. See:

<http://www.caida.org/tools/visualization/geoplot>.

Otter is a tool used for visualizing arbitrary network data that can be expressed as a set of nodes, links, or paths. See

<http://www.caida.org/tools/visualization/otter>.

pchar:

<http://www.employees.org/~bmah/Software/pchar>.

clink: <http://rocky.wellesley.edu/downey/clink/>.

Finally, the graphical user interface will lay out the graph generated by the back end and allow the user to drill down into the collected data and estimated characteristics.

To further strengthen the relation between CAIDA and VU, CAIDA provided us with a GLOBE server in San Diego; in return, VU will set up a skitter machine in Amsterdam. Also, by using CAIDA's data, we can find the best locations for our GLOBE boxes to provide good service coverage, and their data on bandwidth and round-trip time will help us to optimize GLOBE.

Of course the other notable aspects of the exchange were the experiences I had and the new contacts I made. During my work at CAIDA, I met a lot of fantastic people and brought back a number of memories (a rollercoaster park, a real American Thanksgiving, bonfires, parties, etc.) which made the exchange a success.

All in all I found it was a great exchange and would like to thank the following people and institutions for making it all possible: ReX, a USENIX/NLnet endeavor (rex@usenix.org; rex@nlnet.nl); Gale Berkowitz; Frances Brazier; Evi Nemeth; kc claffy; and Maarten van Steen.