
Fathoming AS topology Data

Visualizing Dimes

Mohammed El-Beltagy

mohammed@computer.org



Outline

- Why another visualization tool?
 - The nature of the Dimes Data
 - Mapping AS number to country
 - Data Volatility
 - Demo of the tool
 - Future directions
-

Why another visualization tool?

- None of the visualization tools currently available are well suited to answer questions that are relevant to decision makers
 - How is one country connected to another country?
 - What is the central AS within a particular country?
 - How “scale free” is AS connectivity within a country?
 - How is internet penetration evolving in a particular part of the world?
-

The nature of the Dimes Data

- The released data contains
 - 38,928 AS edges
 - 14,154 AS nodes
 - The topology data is collected over a period of 3 months (3-3-2005 to 31-5-2005)
 - No timestamp on establishing AS links
-

Mapping AS number to country

- We obtain AS number country allocations at any point in time from the five Regional Internet Registries (RIRs)
 - ARIN (<ftp.arin.net/pub/stats/arin/>)
 - RIPE (<ftp.ripe.net/ripe/stats/>)
 - APNIC (<ftp.apnic.net/pub/apnic/stats/apnic/>)
 - LACNIC (<ftp.lacnic.net/pub/stats/lacnic/>)
 - AfriNIC (<ftp.afrinic.net/pub/stats/afrinic/>)
 - Resolve anomalies using the their **whois** database servers
-

Mapping AS number to country

- The combined archived data from the RIRs report 28,901 AS numbers as of 31-5-2005. Comparing this with data at the beginning of DIMES data collection 3-3-2005 we found
 - 1654 “new” AS numbers
 - were not found in the data of 3-3-2005 and were found in the data of 31-5-2005
 - 126 “closed” AS numbers
 - were found in the data of 3-3-2005 and were not found in the data of 31-5-2005
 - 17 “changed” AS numbers
 - were found in both data of 3-3-2005 and 31-5-2005, but the countries that correspond for these AS numbers were changed

Data Volatility

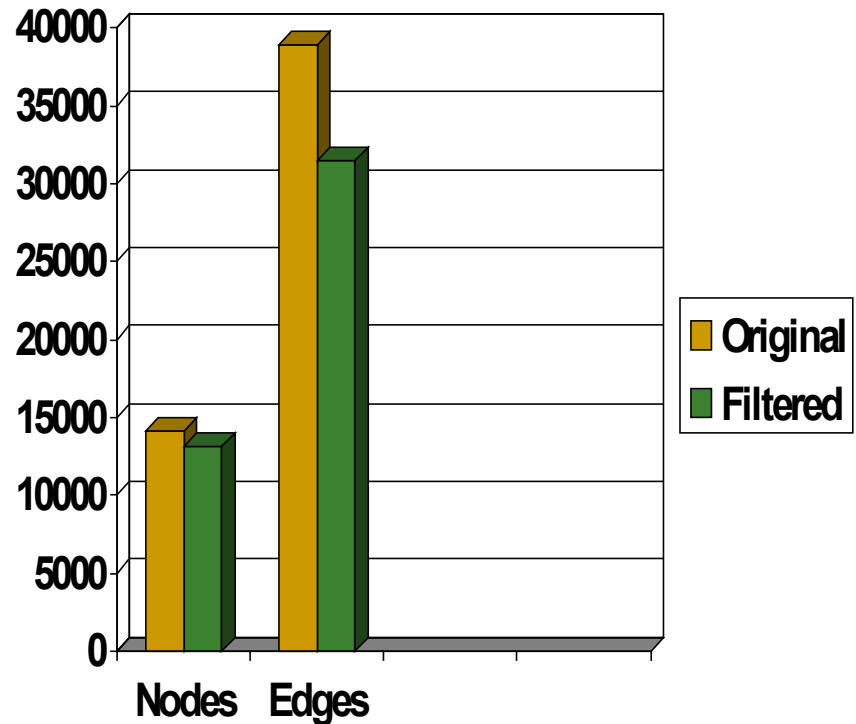
- Using the combined RIR data obtained at 31-5-2005 and trying to map the dimes data we found
 - Mapped Nodes: 13096.
 - Unresolved nodes: 1058.
 - Mapped edges: 31472.
 - Undermined edges: 7424
-

Data Volatility

- Comparing the combined RIR data at the beginning and end of the DIMES data collection to understand what happened to the unresolved nodes we found 12 nodes where **closed** before 31-5-2005
 - That leaves 1046 AS numbers that were most probably allocated and closed within the three month period (7% of the data)
 - Of the 17 **changed** AS numbers only 8 appeared in the DIMES data, so we had to delete 32 associated edges
-

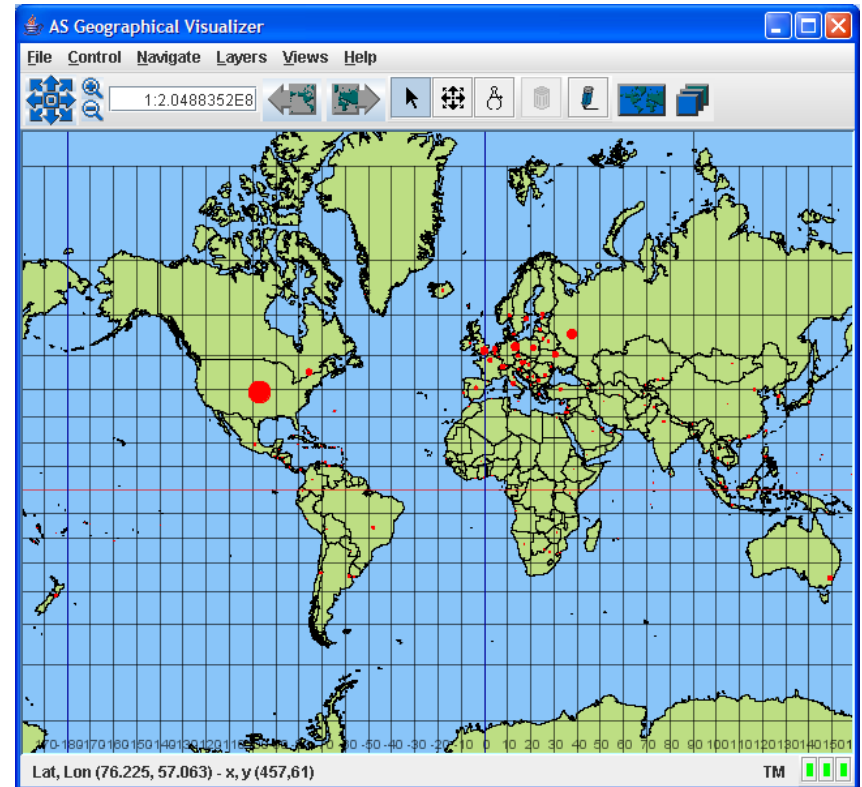
Data Volatility

- We finally have 31,440 edges and 13,088 nodes to work with, out of the original 38,928 and 14,154 nodes.



Demo of the tool...

- We don't know what to call it....



Future directions

- Incorporate Scott's k-core-ness measures
 - Link the data tool with RIR whois servers
 - With time stamped Dimes data we would like to be able to do
 - Temporal queries
 - Understand scale freeness across the world
-