PlanetLab Deployment and Analysis of Network Coordinate Systems

Fenglin Liao
Keshava Subramanya
Veljko Pejovic

{fenglin,keshava,veljko} @ cs.ucsb.edu

Contents

- Motivation for the Project
- Project Objectives
- Experimental Setup
- Challenges
- Results
- Conclusions
- Guidelines for the future
- Credits

Motivation for the project

- Non-availability of exhaustive NC system comparison
- Need for relevant metric for comparing NC systems
- Need for neutral analysis of NC systems under various criteria *

A tale of 2 NC Systems

- Vivaldi:
  - Simple, adaptive, distributed, lightweight
- GNP:
  - Centralized (uses landmarks), static

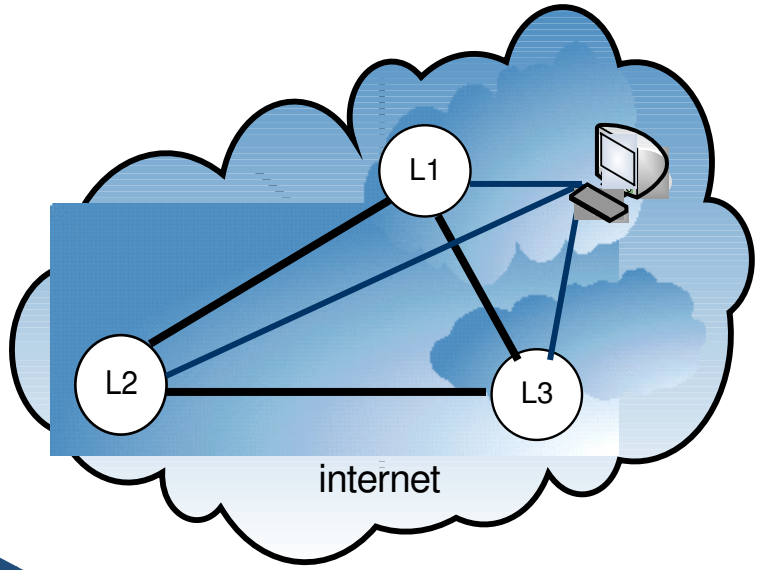

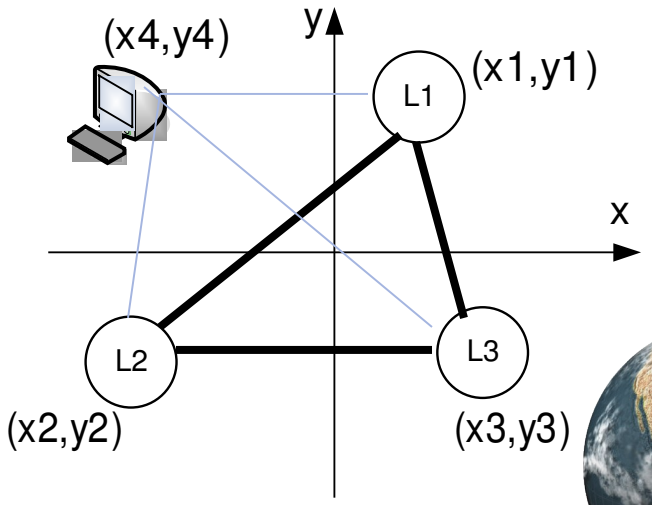

Project Objectives

- Accuracy
  - RTT vs. Calculated distances
  - Effect of node “distance”
  - Convergence time
- Robustness
  - Effect of network traffic
  - Impact of malicious nodes
- Stability
  - Behavior of system when nodes are joining and leaving (Vivaldi)
  - Impact of Landmarks leaving the system (GNP)

Experimental Setup

- Selected 100+ on PlanetLabs
  - Sorted and selected from every class based on RTT
- Ran Pyixda - a Java implementation of Vivaldi based closely on original MIT Paper (Azureus)
- GNP implementation by T. S. Eugene Ng, author of the original paper
- Constructed full ping mesh with periodic ping measurements
- Modified code of Pyxida, created daemons, mined log file data using scripts, analyzed data using Matlab
- Total run time of over 150+ hours

Challenges

- Most importantly, being comprehensive, relevant
- Bringing dissimilar NC systems to comparable platform
- Large amounts of data to be processed
- Need to move massive amounts of Data to PL
  - Killed some ‘csil’ machines in the process
  - Angry emails from local sys-admins due to “ssh probe traffic”

Results - Accuracy

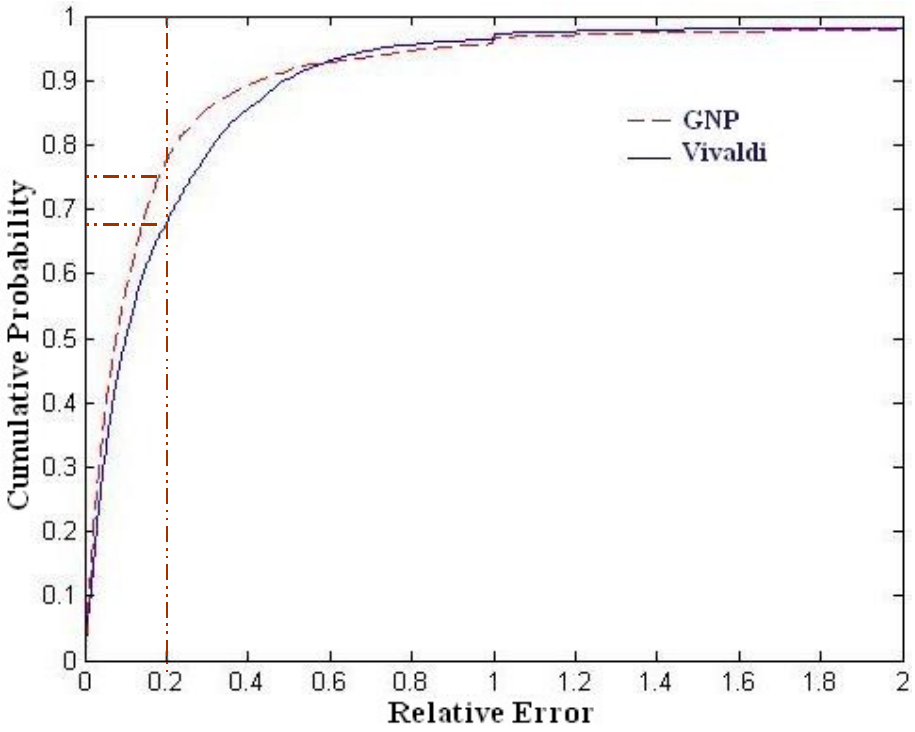

- GNP, as expected, has better accuracy
- Increasing no. of dimensions of Vivaldi improves performance but does not beat GNP
- 65% of Vivaldi, and 75% of GNP measurements have relative error less than 0.2, hence, fairly accurate!

Results - Accuracy: RRL and CNL

- RRL: Relative Rank Loss
  - how well a network coordinate scheme respects the relative ordering of all pairs of neighbors *(Lua et al)*
- CNL: Closest Neighbor Loss
  - Percentage of nodes whose estimated closest neighbor differs from the actual one *(Lua et al)*
- Provides a NC independent platform for analysis of accuracy
- Most applications need just these two, simplifies API set

RRL: Vivaldi vs. GNP

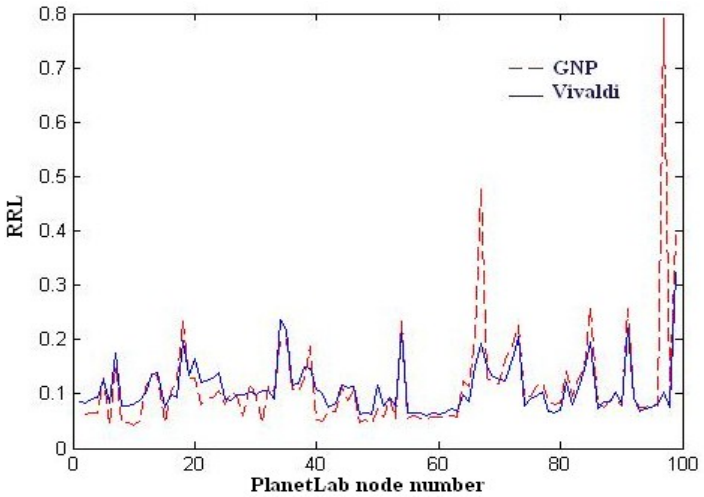

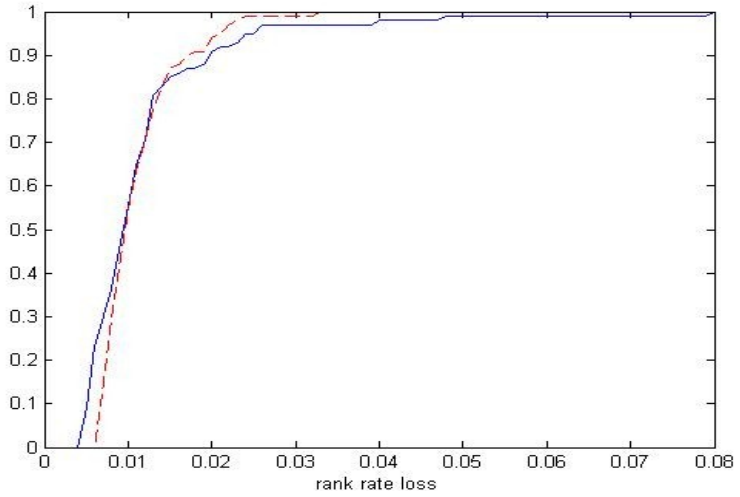

NC	Min. rrl	Avg. rrl	Max. rrl	cnl [%]
GNP	0.0429	0.1154	0.7908	**50.5**
Vivaldi	0.0608	**0.1095**	0.3232	57

- Lower RRL in Vivaldi
- Lower CNL with GNP
- Spikes in GNP for some nodes

Results - Accuracy: Effect of Distance

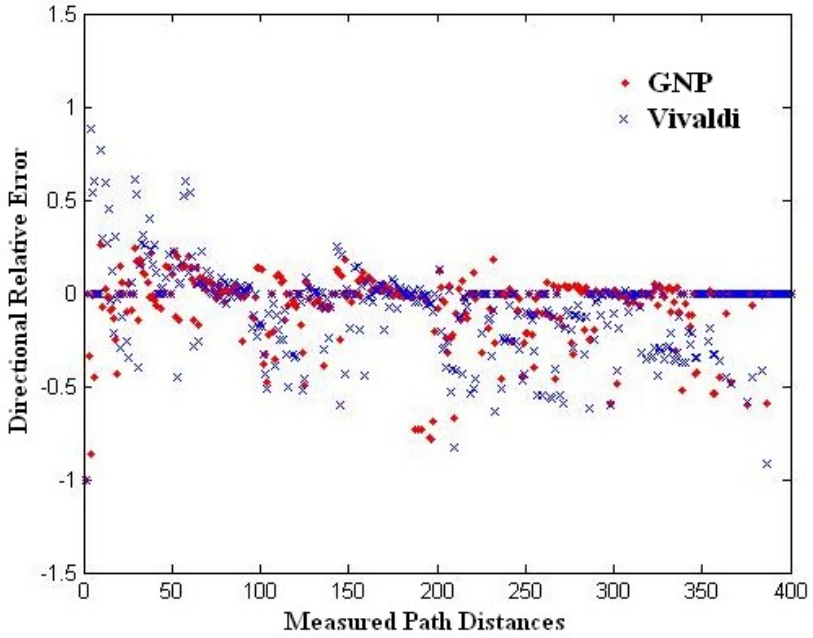

- Vivaldi overestimates closer node distances
- GNP underestimates far away nodes

Results - Accuracy: Convergence time

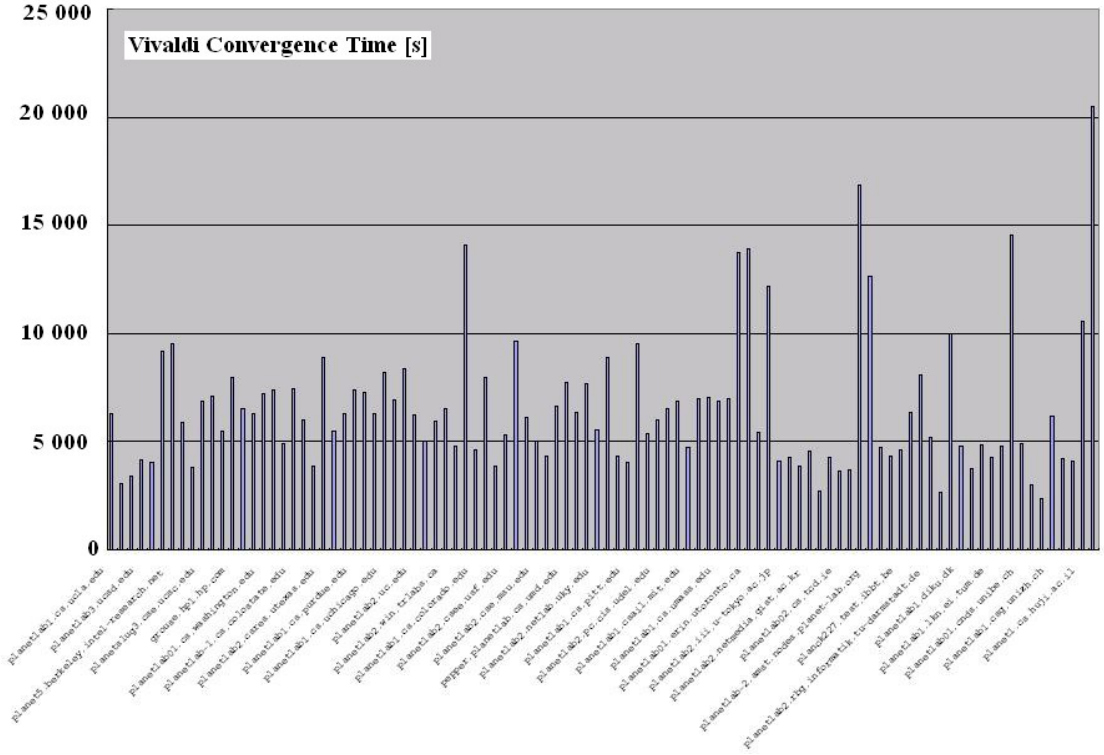

- Most nodes converged (err<0.2) within 2 hours for Vivaldi
- GNP – Static (calculation was CPU intensive)

Results - Robustness: Effect of Network traffic

- Selected 8 nodes and ran Vivaldi and GNP
- Subjected the network to heavy traffic
- Compared measurements before, during and after traffic
- Result:
  - No significant changes for Vivaldi and GNP
  - Possibly because both are really light weight on the network

Results - Robustness: Impact of Malicious Nodes

- Set up a stable ring of 8 nodes
- Hacked the Vivaldi code to send out constant "GossipMessages" with very large values of coordinates
- Added 3 malicious nodes to the network

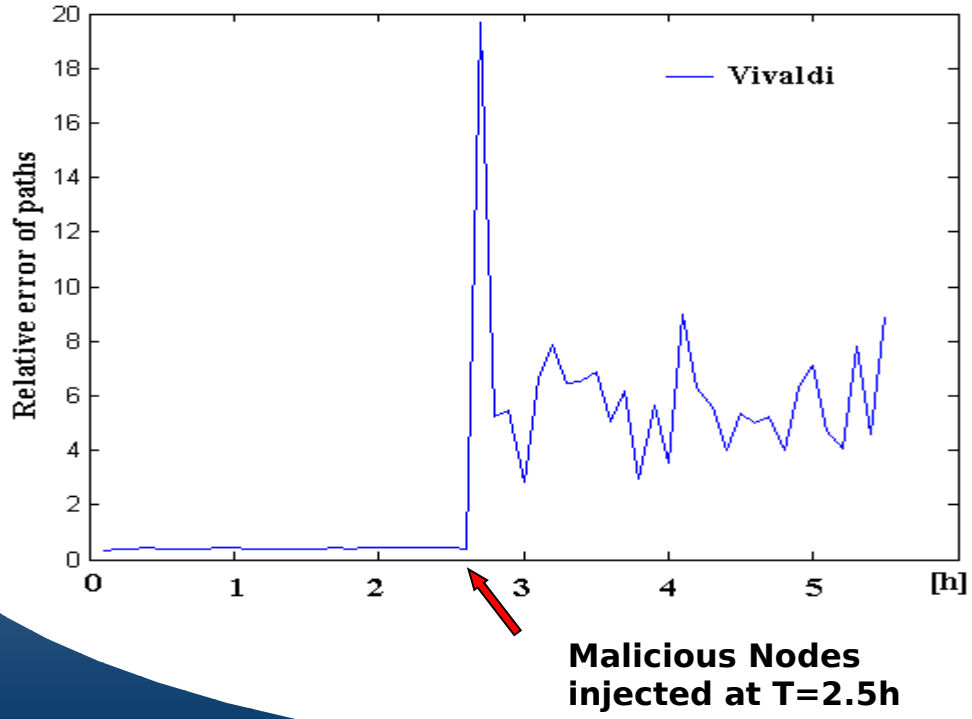

- Vivaldi blown away by the attack
- GNP Malicious Landmark beyond the scope of this work, but Landmarks are usually trusted nodes

Results - Stability: Effect of Network Churn

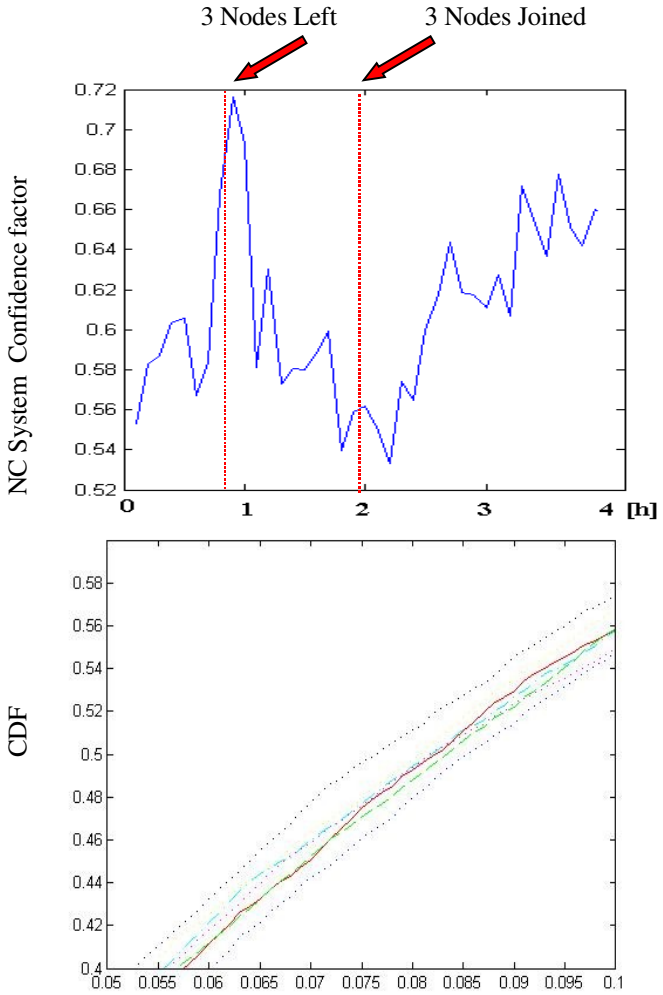

- Vivaldi reacts more to node churning, but recovers soon
- Nodes leaving GNP NC have no effect on other nodes
- Landmarks leaving the network has almost no effect on GNP stability
- “Back-up” landmarks are often available to ensure system stability during landmark failure (called candidate landmarks)
- If the number of available landmarks go below minimum threshold, NC system fails to continue operation

Conclusions

- GNP is generally more accurate but both are still not accurate enough
- CNL performance of both systems are highly inaccurate - more than 50% error
- GNP makes really high RRL on some nodes, hence not very useful for applications relying on RRL
- Based on distance based measurements, Vivaldi is suited for large networks and GNP for smaller and denser networks
- Vivaldi takes about 2.5 hrs to converge. Using it on high churn networks may not be ideal

Guidelines for the future

- A hybrid approach: Vivaldi like, but with periodic random ping to improve accuracy
- Make light weight (a C implementation)
- Multiple implementations (two running - choose one or average estimated RTT)
- Log files to “remember” coordinates of last run (brings down convergence time)
- Vivaldi should guard itself against malicious nodes by using some reputation mechanisms
- Make mobile/wireless nodes “read-only”

Thank You

Questions/Comments?

Credits

- Ledlie, Pietzuch, Parker, http://pyxida.sourceforge.net/
- T. S. E. Ng and H. Zhang. Predicting Internet network distance with coordinates-based approaches
- F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A Decentralized Network Coordinate System. In SIGCOMM, August 2004
- Network Traffic Generator: http://sourceforge.net/projects/traffic
- Harsha Alagud (UCSB) – For his Java expertise and cheerful helping hand