

Demo: Visualization of Stability Monitoring for Node Selection

Thiago Garrett, Luis C. E. Bona, and Elias P. Duarte Jr.

Federal University of Paraná, Dept. Informatics

P.O. Box 19018 Curitiba 81531-980 Brazil

{tgarrett,bona,elias}@inf.ufpr.br

Abstract—The purpose of this demo is to visually show a testbed monitoring strategy used to select “stable” sets of nodes to run new protocols. The stability of a set of nodes is defined in terms of the ability of the nodes to communicate among themselves within given time bounds during reasonable intervals of time. We assume an unstable network, in which some nodes may not be able to communicate with some others, and this condition varies with time. In order to measure stability, the communication between pairs of nodes is continuously monitored by measuring the corresponding Round Trip Time (RTT). A stability graph is generated from the monitoring data in which vertices represent network nodes and an each edge means the corresponding nodes are considered to be stable during an observation period. Multiple different structures have been embedded on the stability graph to select a large enough number of nodes on which the new protocols are executed: based on degree, clique, and k -core. We compare the different strategies both in terms of the quality of the set of nodes returned and how they fare as time passes.

Index Terms—Planetary-scale Networks, Pairwise Monitoring, Testbed Node Selection

I. INTRODUCTION

Before making new network protocols and distributed applications available, it is important to conduct experiments under realistic conditions. Large-scale testbeds have been constructed for 5G networks¹, Internet of Things² and for planetary-scale communications (e.g. PlanetLab³ and M-Lab⁴). However, it is not trivial to select a set of nodes on which to obtain reliable empirical results. Testbeds can be very unstable, in some cases it is important to show that the protocol/distributed application is able to deal with inherent instability. However in other cases this can be a problem, if the set of nodes selected to run an experiment is too unstable it may be even impossible to run some applications.

Researchers often employ simple approaches to select nodes, such as *pinging* all nodes from a single host and selecting those with the lowest RTTs measured from a central location. Other approaches select nodes based on local criteria, such as a recent work [1] proposes the use of fuzzy logic to select nodes of an opportunistic network based on features such as memory and energy available.

¹<https://5ginfire.eu/university-of-bristol-5g-testbed/>

²<https://www.iot-lab.info/>

³<https://www.planet-lab.org/>

⁴<https://www.measurementlab.net/>

In [2] we introduce a strategy to select testbed nodes with a focus on experiment reproducibility. Instead of monitoring the nodes themselves, the proposed strategy samples the RTT obtained from pairwise communications. A stability graph is built in which the vertices correspond to testbed nodes and there is an edge between two vertices if their communication is classified as stable. We investigate the performance of different structures embedded in the stability graph to run experiments on PlanetLab: based on node degree, clique and k -core. Results show that the k -core outperforms the other strategies in terms of their impact on the performance and reproducibility of the experiments.

The purpose of this demo is to visually present results of the stability monitoring strategy, executed on the PlanetLab traces, a very unstable planetary-scale network. The stability of a set of nodes is defined as the ability of these nodes to communicate among themselves within a given time bound during reasonable intervals of time. Different strategies are presented and compared. We also discuss how parameters can impact the graphs and the node selection. Furthermore, in this demo it is possible to visualize how specific sets of nodes selected according to different criteria fare as time passes.

The rest of this paper is organized as follows. In section II we describe the monitoring strategy, and how stability graphs are generated from the acquired data. Section III gives an overview of node selection strategies and show visual examples of these strategies. We draw our conclusions on section IV.

II. MONITORING

In the proposed distributed monitoring strategy pairs of nodes monitor their communication. The basic information obtained is the pairwise RTT. The measurements are done at the application level, in this way measurements are influenced not only by network issues, but also by conditions of the nodes themselves, like CPU utilization for instance.

Samples of the RTT are periodically obtained for each pair of nodes and sent to a central server. This server is responsible for storing and processing the samples and for generating the so-called stability graphs. A stability graph is generated from the monitoring data and represents the stability of the network during a time interval. Each vertex in a stability graph represents a node in the network, and each edge between two

nodes means that at least 90% of the sampled RTTs between those nodes were less than a configurable threshold.

Figure 1 shows an example of a stability graph generated for a period of 1 hour and a threshold of 0.1s. For better visualization, the size of node is proportional to its degree, and nodes are shown organized in clusters of different colors.




Fig. 1: Stability graph for 1 hour period, 0.1s threshold.

III. NODE SELECTION

Node selection is based on the stability graphs. An intuitive approach would be to search for a clique on the stability graph, so that each selected node could communicate stably with all others. However, previous results show that it is unlikely that the same set of nodes will be valid as a clique for a reasonable amount of time [3]. After evaluating different graph structures, we concluded that searching for k -cores in stability graphs is the best approach. A k -core is a subset of nodes that induce a subgraph with minimum degree k , i.e. a subset of nodes that can communicate in a stable fashion with at least k other nodes within the subset.

Figure 2 shows the same stability graph as in Figure 1, highlighting in red nodes selected using the k -core strategy. 41 nodes were selected and the value for k was 35. Note that all selected nodes are in the same cluster.

Figure 3 shows the same stability graph again, but highlighting in red nodes selected according to their degree. Nodes with degree 48 or higher were selected, resulting in 47 nodes. Note that in this case nodes in different clusters were selected. They are well connected within their clusters, thus have a high degree. However, nodes in different clusters are not well connected among themselves. This can result in poor performance when running distributed experiments on those nodes. Furthermore, even nodes in a same cluster may not be connected well enough to run distributed experiments with reasonable performance and repeatability.




Fig. 2: 41 nodes selected with the k -core strategy, $k = 35$.




Fig. 3: 47 nodes have degree at least 48.

IV. CONCLUSIONS

The purpose of this demo is to present visual results of a strategy to monitor and select stable nodes of an unstable network. We described different criteria to select nodes of a stability graph: degree, clique, and k -core. Although the monitoring strategy has been published, the contribution of this demo is to demonstrate the results visually, showing how stability graphs are built, and comparing the different strategies for node selection. The goal is to provide a visual and intuitive demonstration of how the network is monitored and stable sets of nodes are selected.

REFERENCES

- [1] M. Cuka, D. Elmazi, M. Ikeda, K. Matsuo, and L. Barolli, "IoT node selection in Opportunistic Networks: Implementation of fuzzy-based simulation systems and testbed," *Internet of Things*, vol. 8, pp. 100–105, 2019.
- [2] T. Garrett, L. Bona, and E. P. Duarte Jr., "Improving the Performance and Reproducibility of Experiments on Large-Scale Testbeds with k -Cores," *Computer Communications Journal*, vol. 110, pp. 35–47, 2017.
- [3] E. Duarte, T. Garrett, L. Bona, R. Carmo, and A. Züge, "Finding stable cliques of PlanetLab nodes," in *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2010.