

The Practicality of End-User Network Monitoring

Vivek S. Pai
Princeton University

Abstract

With the advent of PlanetLab, the opportunity for the average researcher to monitor a variety of network behaviors from a number of vantage points has increased tremendously. I will briefly discuss the experiences we have had in the following areas: network path anomaly detection in PlanetSeer, detecting anomalous applications in CoMon, and relating our results with those obtained by other groups. Included in the discussion will be where to locate such monitoring, the feasibility of data sharing, and the utility of duplicated effort.

1 Introduction

In a relatively brief span of time, network monitoring has become a very active field of research in the academic community, and by providing a large monitoring infrastructure to almost anyone, PlanetLab has reduced the barrier to entry in this area. The near-term origin of this interest in failure path monitoring is arguably the RON project [1], which demonstrated both practical results and interesting research opportunities. The hurdles toward the “next step” in applying RON-like approaches to PlanetLab, would be in increasing scalability and coverage. Extending an “all pairs” pinging approach does not indefinitely scale, and qualitatively expanding the scope of observation becomes more difficult if the general types of nodes joining the project are similar to existing participants. The diversity of PlanetLab node locations and the paths between them had already been a concern for some research [2].

Within four years, two projects, SOSR [5] and PlanetSeer [8], had taken different approaches to both of these issues. To expand the paths monitored, both systems examined paths outside of PlanetLab. SOSR contacted Web servers using TCP packets, while PlanetSeer used UDP packets to contact Web servers as well as clients using the CoDeeN content distribution network [7]. Both systems also tried to increase the rate of anomalies detected – SOSR performed more active probes, and PlanetSeer passively monitored the existing CoDeeN traffic to determine when to launch active probes.

While both projects were successful in increasing

monitoring scale as well as finding more network anomalies, how to reach the “next steps” are not immediately obvious. Some possible next steps involve how to increase the reach of these systems, how to scale to larger user populations, and how to have more control over the probing. While PlanetSeer did observe nearly one million unique client IP addresses spanning over 9,000 ASes, coverage of Tier 5 ASes is less than 50%. Even in the other AS tiers, where coverage ranges from 80%-100%, coverage is not uniformly distributed, so it is hard to make statements about “the Internet” without global coverage. While CoDeeN’s traffic has doubled to 10-12 million requests/day, and its user population has grown to 50,000 users/day, this is still far short of commercial ISPs.

Even if these numbers grew, the fact that our approaches rely on probing only from the PlanetLab infrastructure means that we may never traverse the actual paths used between clients and servers. One aspect of RON that is necessarily lost in both SOSR and PlanetSeer is the ability to control probing at both endpoints of a connection. PlanetSeer can still monitor some client-initiated traffic, so can determine when some forward paths are failing, but can not determine failures in all forward paths.

2 Going Forward

While actively involving end users in the network monitoring process can overcome some of the limits described above, how to engage them is not obvious. Most of the obvious approaches have privacy-related aspects that could cause significant public-relations problems (or worse) if deployed at scale.

Asking users to actively participate in network monitoring is the most straightforward approach, and the one least likely to cause outrage. However, the number of users willing to participate may not be particularly large, unless some direct benefit is perceived. Consider the Seti@home project [6], which performs distributed signal analysis to search for extra-terrestrial life. It has a total of 5.4 million users, though not all of them may be active at any given time. In contrast, CNN’s Web site has over 20 million unique visitors per month, and AOL has a subscriber count in that same range. If

network monitoring is perceived to be less appealing than searching for alien life, this approach may still fall short of the desired targets.

Using “Web bugs” embedded on the pages of popular sites (directly or via a CDN) may cause large numbers of clients to contact the measurement servers, but this approach may also attract some amount of negative publicity. While this approach does not maintain any control over the client, it has the possibility of generating significant data if popular sites participate. For example, a site’s main page could contain many single-pixel regions linked to measurement servers around the world. If placed correctly, these objects would have minimal impact on download time, and could provide useful data. Standalone pages could even be generated to be used for diagnostic purposes, where the user could load such a locally-stored page when his/her Internet connectivity seems problematic.

However, if we consider the value of the last-mile link, we may arrive at a different approach. Few end-users will be multi-homed, so a system behind a failing last-mile can only record the failure locally and report it when connectivity resumes. As such, ISP-level connectivity may provide nearly the same quality of information as the end-user approach, but with fewer concerns about privacy. We suspect that few ISPs would readily give out their connectivity information, especially if their competitors could use it against them. While this approach is less of a concern for residential customers, it may be an issue for businesses. Instead, if the measurement tools were provided with data access only on a reciprocal basis, the ISPs would have some reason to run them. Embedded this data into a necessary protocol, such as BGP, would be a more ambitious future step, but might also be possible.

In the near term, it is likely that many of these approaches will be implemented, and that some of the effort involved will be seemingly duplicated. Over time, whichever efforts gain the most traction will undoubtedly gain more attention, but this process is arguably better than attempting to choose *a priori* the approach we expect to win. As long as some mechanism exists for collecting the data from any source, extra coverage is only a minor waste of bandwidth.

We have two experiences in this kind of duplicated measurement, and both have been positive. It is interesting to note that SOSR and PlanetSeer reached seemingly different conclusions while performing similar studies, but a closer examination reveals why this disparity is not real. By examining data at the path level, SOSR concludes that most paths function well. However, since PlanetSeer is concerned with the aggregate quality of the CoDeeN service, observing two network failures per minute appears alarming. How-

ever, if one were to translate the PlanetSeer numbers into per-path failures, the data looks similar to that observed in SOSR.

The other case centers around measurement of PlanetLab itself, and tries to detect anomalous applications or nodes. Our group developed CoMon [4], which aggregates node-centric and application-centric data on PlanetLab. In parallel, a group at Intel developed PSEPR [3] which performs more correctness monitoring. CoMon and related tools incorporate some data from PSEPR, and in some cases, this data differs from what their own measurements indicate. In these cases, the differences have been isolation-related – the two tools, despite running on the same machine and measuring similar quantities, observe different outcomes due to some bug in the environment, generally related to the mechanisms used to isolate applications.

In conclusion, many reasonable paths for the “next step” in large-scale network measuring exist, and some combination of them will likely yield more data than we can harness now. I expect that in a relatively short period of time, these approaches will provide a standard basis for the next advances in measurement, and will begin to provide a much more complete picture of network paths and their failures.

Acknowledgments

Many of these ideas emerged and have evolved over time in discussions with Larry Peterson, Ming Zhang, and others.

References

- [1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *Proc. 18th ACM SOSP*, Banff, Canada, Oct 2001.
- [2] S. Banerjee, T. G. Griffin, and M. Pias. The interdomain connectivity of PlanetLab nodes. In *Passive and Active Measurement Workshop*, April 2004.
- [3] P. Brett, R. Knauerhase, M. Bowman, R. Adams, A. Nataraj, J. Sedayao, and M. Spindel. A shared global event propagation system to enable next generation distributed services. In *Proc. WORLDS 2004*, Dec 2004.
- [4] CoMon. <http://comon.cs.princeton.edu>.
- [5] K. P. Gummadi, H. V. Madhyastha, S. D. Gribble, H. M. Levy, and D. Wetherall. Improving the reliability of internet paths with one-hop source routing. In *Proc. 6th OSDI*, San Francisco, CA, Dec 2004.
- [6] Seti@home. <http://setiathome.ssl.berkeley.edu/totals.html>.
- [7] L. Wang, K. Park, R. Pang, V. S. Pai, and L. Peterson. Reliability and Security in the CoDeeN Content Distribution Network. In *Proc. USENIX ATC 2004*, Boston, MA, June 2004.
- [8] M. Zhang, C. Zhang, V. Pai, L. Peterson, and R. Wang. PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services. In *Proc. 6th OSDI*, San Francisco, CA, Dec 2004.