DIMES

Letting the Internet Measure Itself

Yuval Shavitt & Eran Shir
Tel-Aviv University
www.netDimes.org

Yuval Shavitt's email: shavitt@eng.tau.ac.il, Skype: yshavitt
Eran Shir's email: shire@eng.tau.ac.il, Skype: theshire

Abstract

Today's Internet maps, which are all collected from a small number of vantage points, are falling short of being accurate. We suggest here a paradigm shift for this task. DIMES (Distributed Internet MEasurement & Simulation) is a distributed measurement system for the Internet that is based on the deployment of thousands of light-weight measurement agents around the globe. We describe the rationale behind its deployment and present initial analysis of its results in the first six months of operation. We demonstrate its effectiveness compared to maps created using BGP (Internet routing protocol) update reports.

Introduction

As the Internet evolved rapidly in the last decade, so has the interest in measuring and studying its structure. Numerous research projects [1], [2], [3], [4], [5], [6], [7], [8] have ventured to capture the Internet's growing topology as well as other facets such as delay and bandwidth distributions, with varying levels of success. As the Internet continues to grow, especially far from its North American based core, measurement discrepancies are growing as well. A main handicap of current measurement projects is their rather limited number of measurement nodes (usually a few dozens up to a few hundreds) causing results to exhibit bias towards the core. In order to remedy this situation, a measurement infrastructure must grow several orders of magnitude in size and global dispersion.

We present DIMES, a highly distributed, global Internet measurement infrastructure, with the aim of measuring the structure and evolution of the Internet using a large set of interacting measurement agents. The key shift suggested in DIMES is the moving from a small set of dedicated nodes, with measurements as their virtually sole objective, to a large community of host nodes, running light weight low signature measurement agents as a background process. Given the importance of location diversity in Internet measurements, this shift promises to enhance measurement results considerably.

Figure-1 **Agent-Server Interaction Design**

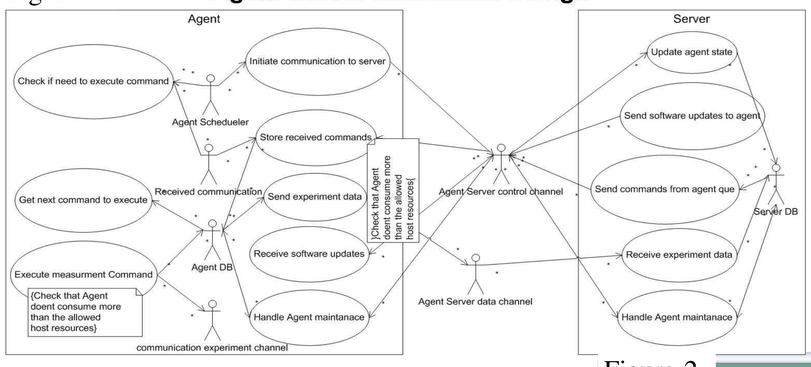

DIMES Measurement Goals

1. Measure the entire accessible AS graph.
2. Gain agent presence in 95% of ASes in the Internet in three years.
3. Create a fresh AS-level graph every two hours in eighteen months.
4. Create a fresh Router-level graph every week in eighteen months.
5. Annotate the Internet graphs with available bandwidth information in twelve months.
6. Measure the mobile/cellular Internet using a mobile DIMES agent in two years.

DIMES Analysis Goals

1. Create dynamic models which reproduce the Internet's evolution as portrayed in the DIMES measurement results, both for the AS (organizational) level and for the Router level.
2. Infer an accurate PoP (Point of Presence) level graph.
3. Develop social/economic growth metrics based on Internet evolution.

Architecture

The DIMES architecture is a simple multi-channel client-server architecture, based as much as possible on COTS technologies (such as MySQL DB, Tomcat App. server, Apache web server). The agent is written mostly in Java and partly in C++. As can be seen in figure 1, the two communication channels are the control channel (which is secured and authenticated) and the data channel (which is a clear channel).

The main design goals implemented in the DIMES platform were:

a. Agents Security - achieved by using an encrypted, secure channel for all control communication, especially software updates, making it very hard to hijack a DIMES agent. In order to assure user concerns, the DIMES code is distributed under an open source license.

b. Maximum flexibility in data manipulation - achieved by using a robust relational database (MySQL) for data storage and manipulation. This decision is probably the most fruitful design decision we have made.

c. Enabling multiple experiments - achieved by implementing an asynchronous experiment management system which allows the provision of multiple scripts and their execution according to priorities.

Other features that we implement are: a. rate limiter in the agent, which controls the agent's bandwidth usage so that the agent won't be considered a hacking tool by various net admins. b. a rich scripting language which enables sophisticated (timed) experiments. c. a plug-in mechanism for the measurement modules, allowing the introduction of new measurement types.

In order to provide incentives for volunteers to install our agent, we are working hard on making it very useful and catching. As can be seen in figure 2, we provide users with a graphical view of their agent's measurements, visualizing IP links and annotating the graph with IP affiliation data such as country and AS. The users can also perform their own traceroutes, thus making the DIMES agent one of the most intelligent traceroute clients existing.

The first measurement modules deployed in the agent are Ping and Traceroute:

Ping - There are two ping types Implemented in the agent, one which is ICMP_ECHO based and one which is UDP based. Two types are needed in order to circumvent, where possible, various firewalls and filters. In addition, we find it useful when attempting to perform router resolution by comparing destination address and returning address.

Traceroute - Our traceroute (which also has two flavors) is ran several times in each instance, cleaning delay artifacts. Since even between two consecutive traceroute the path may change (for instance due to load balancing), Our traceroute returns a weighted DAG rather than only the first or the dominating path.

Figure-2

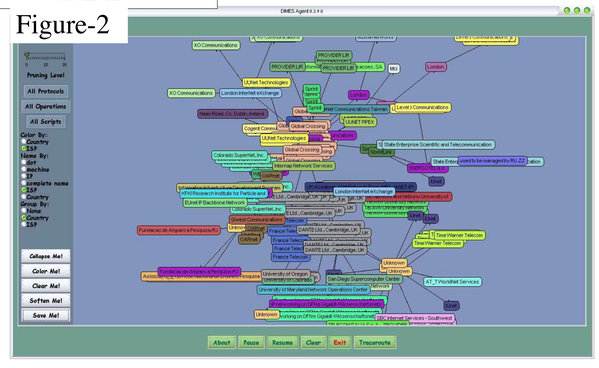

DIMES Results and Analysis

Methodology - AS Inference

IP-level hops are identified from the raw data received from the agents. The AS affiliation of IP addresses is inferred by performing a longest prefix matching search using an up-to-date prefixes database gathered with an IRR server. The AS inference of an edge is performed only if the two nodes of the edge have responded. Thus, if one of the nodes is unknown, the path resolving skips to the next hop which consists of two answering nodes.

Our analysis is performed on the resulting undirected graph which is the integration of measurement taken over the six months since initiation.

Methodology - BGP Comparison

DIMES results are compared to the state of the art topology aggregation resource, which is the BGP updates gathered in RouteViews [9]. In order to compare on a common ground, the BGP graph is created by gathering BGP updates from every day during the same six months. Again, we analyze the resulting undirected AS graph.

DIMES-BGP comparison

As can be seen in figure 3 below, though DIMES still covers less ASes than BGP data, it finds many edges (11,300 to date) that are unknown to BGP, even when integrated over six months. This makes the complete Internet graph BGP and DIMES unification) much denser than previously reported, especially far from the core. This result is further illustrated in figure 4 which compares AS degrees in the BGP and the complete graph.

Figure-3 AS map on March 1st

Figure-4

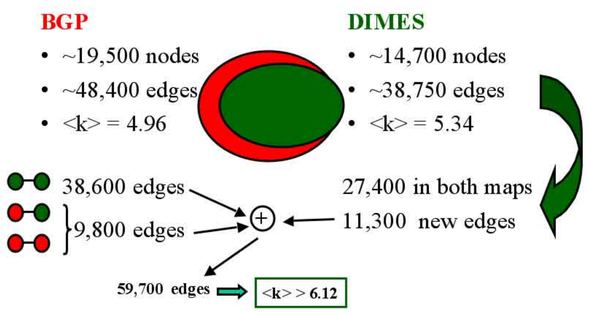

Figure-5

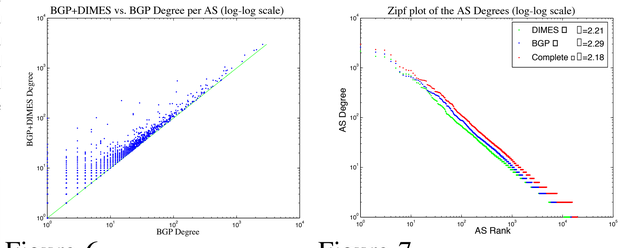

Figure-6

Figure-7

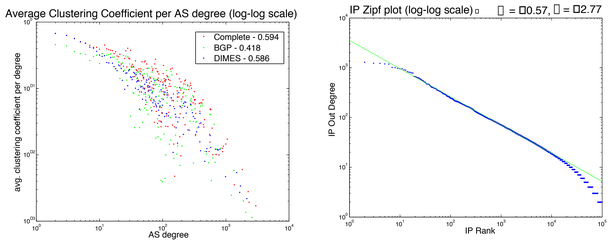

Figure-8

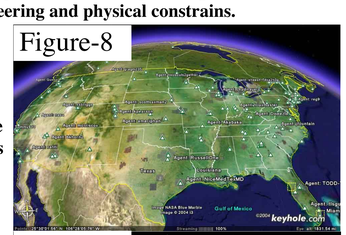

Figure-9

Figure-10

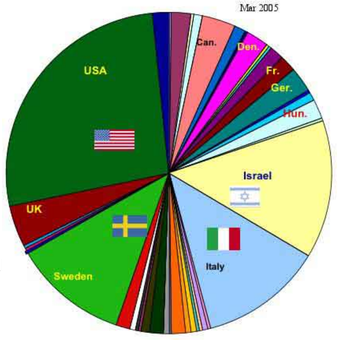

Internet Graph Properties

The AS graph degree distribution exhibits a power law as expected, both for the BGP data set, the DIMES data set and their unification (Complete). As can be seen in figure 5, the respective Zipf plots are very close, and the power exponents are within a 10% range from each-other.

While having a similar degree distribution, the DIMES and BGP graphs are dissimilar in many other facets, the most striking of which is their clustering coefficient and clustering distribution. As presented in figure 6, the BGP graph's clustering coefficient is only 0.418 while DIMES' coefficient is 0.586 - a striking 40% increase. In figure 6 one can also observe the difference in the clustering distribution, especially in the low degree nodes. This observation illustrates the inherent bias towards the core that BGP based topologies suffer from. It shows that the periphery of the Internet is much more inter-connected than what was assumed.

In figure 7 we present the Zipf plot for the IP-level graph degrees which exhibits a very smooth power distribution over all routers with minimal degree between 10 and 1000. Its power exponent is higher than its AS counterpart, as can be expected, given that unlike in the AS case, it is heavily dependent on engineering and physical constrains.

DIMES Community

DIMES' success heavily depends on our ability to harness a large, dispersed community of volunteers. Currently (March 1st.) the DIMES community consists of more than 800 agents from more than 50 countries performing together about 150K measurements per day. Figures 8-9 present the geographical dispersion of agents in north america and the US. In addition we have considerable presence in South-East Asia and in Australia. Figure 10 presents the agents distribution by country.

DIMES Future

Our ultimate goal for DIMES is to fill with content the phrase "Let the Network Measure Itself". We envision an Internet in which, just like in our body, each element takes part in the monitoring and regulation task. In the near future we will reach a point where DIMES will produce a fresh map of the Internet every two hours or less. Such resolutions will allow it to become an operational tool helping in maintaining the Internet's healthcare.

In addition to developing new correlated measurement types we aim to move DIMES to the mobile realm with mobile and embedded DIMES agents. We aim to create a global Internet measurement resource for researchers, which will be able to use our data, perform their own experiments, and even suggest new types of measurements to be deployed in the agent. If you are interested in one of these aspects you are welcomed to contact us.

Acknowledgments

We would like to thank Anat Halpern and Ohad Serfaty for their wonderful work, helping make DIMES a reality.

ES would like to thank Prof. Sorin Solomon for his insights and guidance.

The DIMES project have been supported by the EU 6th FP, IST Priority, Proactive Initiative "Complex Systems Research", as part of the EVERGROW integrated project.

ES was partially supported by the "Yeshaya Horowitz Association through the Center for Complexity Science"

Last but not least, we would like to thank all of the volunteers who help us perform this research, and welcome you to join them at: www.netDimes.org

References: [1] M. Faloutsos, P. Faloutsos, and C. Faloutsos, ACM SIGCOMM 1999, Boston, MA. [2] L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos, Global Internet, Nov. 2001. [3] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, IEEE Infocom 2002, New-York, NY, USA, Apr. 2002. [4] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie, IEEE INFOCOM '03, San Francisco, CA, USA, Apr. 2003. [5] S. Bar, M. Gonen, and A. Wool, PAM '04, Antibes Juan-les-Pins, France, Apr. 2004. [6] P. Barford, A. Bestavros, J. Byers, and M. Crovella, ACM SIGCOMM IMW'01, San Francisco, CA, USA, Nov. 2001. [7] R. Govindan and H. Tangmunarunkl, IEEE Infocom 2000, Tel-Aviv, Israel, Mar. 2000, pp. 1371–1380. [8] A. Broido and K. Claffy, SPIE International symposium on Convergence of IT and Communication'01, Denver, CO, USA, Aug. 2001. [9] "University of Oregon Route Views Project," http://www.routeviews.org.