EVERGROW
www.evergrow.org

Project no.	001935
Project acronym:	EVERGROW
Project title:	*ever-growing global, scale-free networks, their provisioning, repair and unique functions*

INTEGRATED PROJECT
IST PRIORITY

Deliverable reference number and title:	D.2c.3 Analysis of Extrapolations from EVERGROW data
Due date of deliverable:	December 31, 2007
Actual submission date:	December 31, 2007

Start date of project:	January 1, 2004
Duration:	48 months
Organisation name of lead contractor for this deliverable:	TAU
Revision:	1
Dissemination level:	PU

Contents

1 Executive summary

DIMES is a large-scale distributed measurements effort that measures and tracks the evolution of the Internet from hundreds of different view-points. DIMES collects, processes and archives measurements data in order to provide a better understanding of the Internet topology and evolvement. As such we have spent much effort in 2007 on making DIMES capabilities available to researchers at several levels of detail and involvement.

An important aspect of DIMES is understanding how researchers should use the data, and to be able to quantify and minimize the intrinsic bias that might exists in the data collected and processed. This deliverable provides an in-depth analysis of the bias in DIMES and methods to minimize it, so that further analysis will be performed on "clean" data.

2 Milestones in the Deliverable

Milestone M.2c-2.4 is in part accomplished by the work done across the EVERGROW project and reported in Deliverable D.1c.2. In this deliverable, we address the issues of data quality and possible bias that present possible limits to reliable extrapolation and modelling when based on the data captured by the DIMES distributed topology mapping effort.

3 Contractors contributing to the Deliverable

This deliverable was done in TAU in collaboration with HUJI and in consultation with other members of EVERGROW, i.e., COLBUD, SICS and KTH.

4 Results

DIMES is a large-scale distributed measurements effort that measures and tracks the evolution of the Internet from hundreds of different view-points, in an attempt to overcome the "law of diminishing returns" [6]. In weekly average (of recent weeks), DIMES has over a thousand active agents, that contribute over 5 million weekly measurements from over 150 different vantage points, geographically distributed around the globe.

In order to create AS-level topology from the IP-level traceroutes provided by DIMES agents, AS resolution is performed for each hop in all paths. Although IP-to-AS mapping can be a difficult task [9, 8] we take a straight-forward resolution approach, and then filter out links that are suspected to be incorrect. AS resolution is done by first performing longest-prefix-matching against BGP tables obtained from RV archive. This resolves approximately 98% of the IP addresses. For the remaining 2%, we query against two WhoIs databases, namely RIPE and RADB, that resolves additional 1.5% of the IP addresses. The remaining 0.5% unresolved IP addresses are discarded and do not participate in the inference algorithm.

Each of the DIMES agents perform measurements by following a script that is sent to it from a central server. The agent can perform Traceroute and Ping measurements using ICMP or UDP packets. The agents periodically access the central server and request a new measurements script to perform.

The server has the ability to send two types of scripts. The first type are scripts that are planned by researchers. These script follow some logic that can be planned by researchers and added to the server. When an agent accesses the server, the server checks if there is a planned script that is waiting for that particular agent, and if so, it will send the script.

If the server does not have a script for an agent, it will give the agent a random "default" script. The default scripts cover the entire IP prefix space in a random manner. We collect the list of prefixes from the RouteViews BGP tables. Using each prefix we construct IP destinations by finding three host IPs in the class of the prefix (we use 1,2 and 254 as the host numbers in order to construct an IP address). A default script includes traceroute and ping commands to 50 destination IPs.

4.1 Filtering Data

The raw DIMES data is filtered in order to reduce inference mistakes and inclusion of false links. First, all traces that exhibit some known traceroute problems [7], namely routing loops and destination impersonation, are detected. Within these tracees, only the section of the path preceding the identified problem is used for topology inference.

There are several methods that can be used for furthering filtering of the data. The first method is to include edges only if they are seen more than once. However, this method might not be accurate since a single agent might report more than once on the same non-existing edge.

A complementary filtering method, is to include edges that are seen by at least two different agents. This, however, has two undesired effects. The first is that this can cause many important edges to be dropped from the topology, since there are agents in remote locations that measure periphery edges, unseen by any other agent. The second is when users form "groups" of several agents, sometimes co-located in a small geographical section. In this case, unwanted edges that are incorrectly discovered by the group of agents will be included in the topology. This filtering method is far from providing a good enough result.

In order to avoid some of these difficulties, a different filtering process is employed. First, the discovery of the AS (or ASes) that are the source of the measurements for each of the measuring agents is performed. The discovery is done by following each of the traceroutes of a given agent, until reaching a hop with a routable IP address that can be resolved into a valid AS. If we find such an IP in the first 4 hops in the path, this AS is considered to be the Origin AS of the agent. A simple yet quite effective filtering method can thus be to include only edges that are seen from more than one origin AS. However, this method also has its flaws.

Some routers incorrectly respond to ICMP queries as if they were the destination, and return a TTL_EXCEEDED message with the source IP address of the destination instead of their own. If such a router is located near the agent (less than 4 hops away), it causes the agent to have many different origin ASes and might lead to non-existing edges.

In order to avoid this, the number of measurements that an agent performs from each of its origin ASes is calculated. The Dominant Origin ASes are the ASes that are the origin of least 50 measurements. Only edges that were discovered using traces carried from the dominant origin ASes are included in the topology. Since most of the agents are not mobile, at most 4 dominant origin ASes are discovered for all agents.

In order to be sure that an AS edge really exists, the unique set of origin ASes is collected (referred to as the *Vantage Points* set). Only edges that are seen from more than a single Vantage Point are included in the topology.

This filtering process result in a very clean and robust AS-level topology, that can be used in conjunction with other data sources in order to create a more complete topological view of the Internet. In Section. 5.1 we evaluate the filtering methods proposed using other sources of topological data.

5 Measurements Bias Analysis

5.1 Filtering Topology

As discussed in Section. 4.1, several methods of filtering can be employed on data derived from DIMES in order to obtain a "clean" AS-level topology. Table. 1 shows the resulted topologies using data from week 50 of 2007, with different methods of filtering employed. The different DIMES topologies are correlated with the topology inferred using RouteViews BGP data.

The table shows that accurate vantage point filtering provide a topology that is very similar to the one without filtering at all, while the other filtering techniques remove much more edges from the topology. This might suggest that VP filtering is much finer than the other methods, and it manages to carefully take out unwanted edges. However, while it is quite difficult to know which method is the right one, it

Table 1: Filtering effect on topology

Filter	Total Topology		DIMES Topology		Edges Origin		
	ASes	Edges	ASes	Edges	DIMES	RV	Both
None	27284	76372	19068	51633	19578	24739	32055
Dates	27277	66815	18378	39343	10021	27472	29322
ASes	27279	68997	18708	42478	12203	26519	30275
VPs	27283	75935	18733	43790	19141	24745	32049

is clear that one can choose a method that suits his conservatism best.

5.2 Vantage Points Stability

Since DIMES completely relies on users to run agents, it is important to validate the stability of vantage points over a period of time. One should expect that not only there are more than 250 vantage points, but these vantage points stay consistent over time. Fig. 1 shows the vantage points DIMES agents measured from over a period of 6 weeks in 2007. The vantage points AS number is translated to an index number. The figure shows that most of the vantage points remain constant over the measured period.

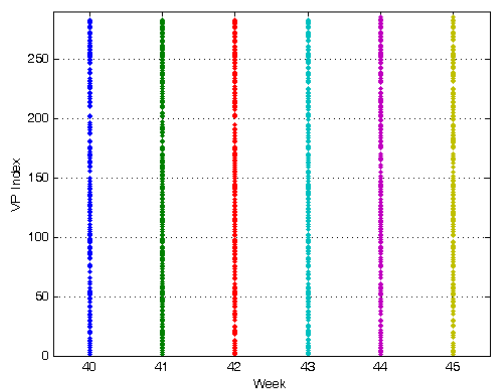

Figure 1: Stability of vantage points

Table 2 provides the number of vantage points for each of the weeks plotted in Fig. 1. It shows that there is large variance in the number of measurements and vantage points when comparing different weeks. Although it seems like a descending trend, looking at the same data from the first weeks of 2008, we have an average of over 200 vantage points and more than 4 billion weekly measurements.

5.3 Agents and Measurements

Each vantage point can hold several agents, that together contribute ASes and edges to the topology. Fig. 2 shows the number of agents that co-exist in an AS per a given AS degree. It shows that while there are many vantage points that have a single agent in them (mostly mid-level degrees), there are many vantage points, in a wide variety of degrees, that hold many agents.

Table 2: Number of vantage points and measurements for weeks 40 till 46 of 2007

Week	Vantage Points	Measurements
40	155	6569737
41	206	14155461
42	179	9799113
43	185	6327969
44	173	6942287
45	165	5737947
46	215	4033571

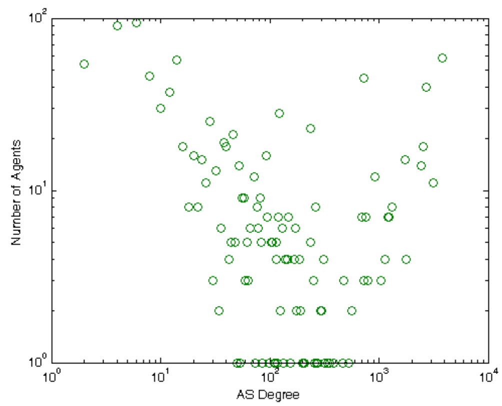

Figure 2: Number of agents by vantage point degree

Fig. 3 plots on a log-log scale, the number of measurements that are performed from per degree of the vantage points. It shows that there is no direct relationship between the number of measurements to the vantage point degree, since it is possible to see low-degree vantage points that preform much more measurements than high-degree vantage points. This shows that there is a good distribution of measuring agents across the different vantage points.

5.4 Diminishing Returns

Barford *et al.* [12] showed that the utility of adding vantage points beyond the second one, quickly diminishes in terms of nodes and edges. In order to examine this claim, we calculated the number of edges and ASes discovered by each of the vantage points. Fig. 4 shows the effect that adding vantage points (sorted by a non-increasing order of the number of measurements performed from each) has on the number of ASes and edges that exist in the topology. It shows that both the number of ASes and especially the number of edges gradually increase even when reaching over 100 vantage points.

This proves that although the diminishing return effect exists, it takes much more than a few vantage points until it “kicks-in”. Had we sorted the vantage points in other methods, this effect could have been even more noticeable. This is mainly due to the fact that lateral edges, i.e., links between ASes in the edges of the Internet, that participate in paths that do not always traverse the core of

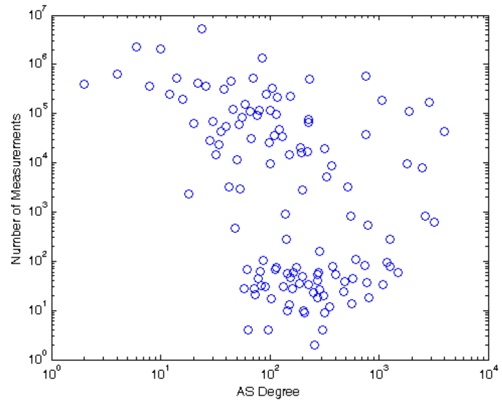

Figure 3: Number of measurements by vantage point degree

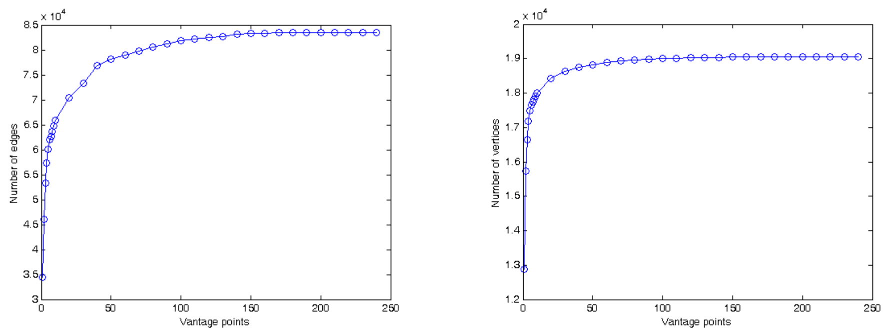

(a) Edges

(b) ASes

Figure 4: Diminishing return

the Internet, are extremely hard to find, and require a very high distribution of the measuring agents.

5.5 Degree Bias

Lakhina *et al.* [13] showed that AS degrees inferred using traceroute-like sampling technique, are highly affected from the location of the vantage points that perform the measurement and collection of data. The authors claimed that this bias could cause the power-law distribution, and suggested two criteria for detecting bias in traceroute studies: a) do the highest-degree ASes tend to be near the measuring sources? and b) is the distribution shape near the measuring sources different from that further from the source?. Using these criteria, the authors showed some commonly used datasets, including Skitter exhibit bias.

Fig. 5 plots the average distance of all ASes with a given degree to the nearest vantage point. It shows that while it is possible to see a correlation between the distance to small and high degree ASes, the variance is rather high, and the distribution of distances of mid-level degree ASes is quite uniform.

Calculating the distance from nearest hops histogram showed that most ASes are 1 to 3 hops away from the nearest vantage points, and a small fraction of the ASes have either a vantage point that exists in them, or have vantage points that

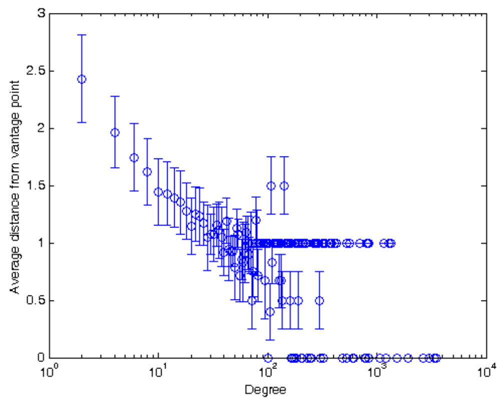

Figure 5: Average distance from nearest vantage points

are 4 or more hops away. This shows that the agents in DIMES are well-spread and provide a good coverage of the Internet at the AS-level.

5.6 Clustering and Betweenness

Some other factors that we wish to further investigate in a planned paper are the clustering coefficients (CC) and betweenness centrality (BC) of the DIMES data comparing to other well-known databases. Fig. 6 provides a comparison of CC and BC of the RouteViews BGP data, DIMES data, CAIDA's Skitter data and merged topology of all. Both plots show that DIMES exhibits the closest match to the merged topology, despite the fact that DIMES and RouteViews are relatively equal contributors of edges to this topology. This might suggest that despite missing AS edges, DIMES does manage to capture rather accurately the properties of the AS-level Internet topology. We plan to publish a full analysis of this data in an upcoming paper.

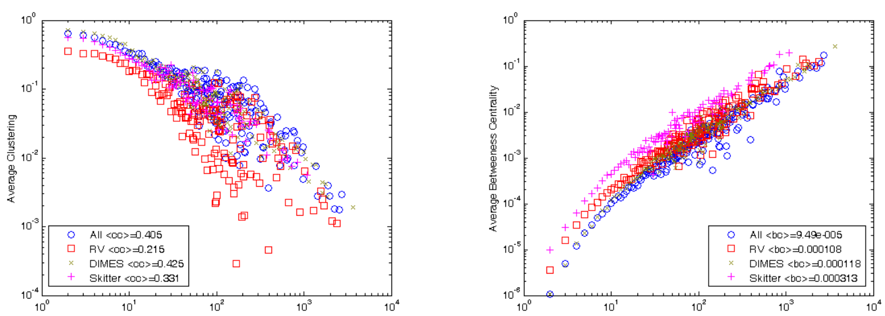

(a) Clustering Coefficient

(b) Betweenness Centrality

Figure 6: Comparison of the Clustering Coefficient and Betweenness Centrality of DIMES, RouteViews, Skitter, and a merged AS-level topology

5.7 Visibility of AS-links

Another interesting property we wish to further examine is the difference between the visibility of AS-links between different topology measurement efforts (i.e., Route-Views, Skitter and DIMES). The most obvious expected difference is that RV uses passive collection of BGP announcements while Skitter and DIMES use active measurement. However, Skitter and DIMES also differ; while Skitter uses few but mostly-active monitors, DIMES uses thousands of agents that are not always active.

These differences result in a different visibility distribution of edges across data collected over a week. Fig. 7 shows the distribution of days (collected during week 41 of 2007) that AS-edges where visible in each of the data sources. It clearly shows that almost all RV edges are seen throughout the entire week. Skitter and DIMES distribution is rather similar, although DIMES exhibits slightly more shortly-visible AS-edges. This is mainly the result of agents that are turned off some days of the week.

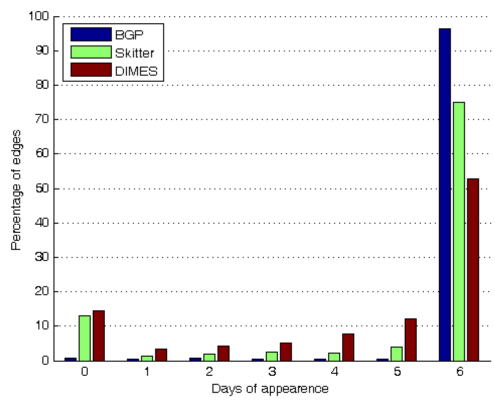

Figure 7: Visibility of AS Edges

Looking at the distribution of the exact day-in-week of DIMES edges that appeared for a single day, did not reveal any interesting finding. This suggests that there is no specific day on which very short-seen edges are visible in the DIMES system.

5.8 Degree and k-Shell Decomposition

Carmi *et al.* [17] presented the *Medusa* model, that uses a k-pruning algorithm to decompose the Internet AS graph and extract a nucleus (the K_{max}-Core) which is a very well connected globally distributed subgraph. This algorithm extracts a core by looking at the entire graph.

Fig. 8 plots the degree of ASes in each kShell on a semi-log scale. We plan to study the difference between these distribution, and understand the effect they have on the resulting topology.

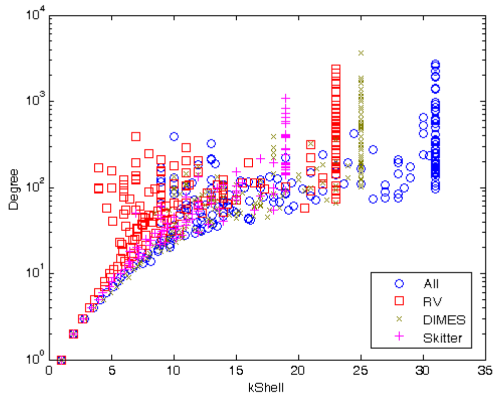

Figure 8: kShell vs. Degree distribution

6 Other material

Not Applicable

7 Papers and publications

References

[1] Adam Shaked Gish, Yuval Shavitt, and Tomer Tankel. Geographical Statistics and Characteristics of P2P Query Strings *The 6th International Workshop on Peer-to-Peer Systems (IPTPS'07)*, Bellevue, WA, USA, February 2007.

[2] Yuval Shavitt and Eran Shir. DIMES: Let the Internet Measure Itself. In Michael H.W. Weber, editor, *"Distributed & Grid Computing - Science Made Transparent for Everyone. Principles, Applications and Supporting Communities"*. Tectum Verlag, to appear Q1, 2008.

[3] Eli Brosh, Asaf Levin, and Yuval Shavitt. Approximation and Heuristic Algorithms for Minimum-Delay Application Layer Multicast Trees. *IEEE/ACM Transactions on Networking, 15(2): 473–484, April 2007*, accepted.

[4] Yuval Shavitt and Tomer Tankel. On Internet Embedding in Hyperbolic Spaces for Overlay Construction and Distance Estimation. *IEEE/ACM Transactions on Networking, 16(1), February 2008*, accepted

[5] Y. Shavitt, E. Shir, and U. Weinsberg Near-Deterministic Inference of AS Relationships In arXiv:0711.4562v1, Nov. 2007.

[6] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *IMW'01: the 1st ACM SIGCOMM Workshop on Internet Measurement*, pages 5–17, 2001.

[7] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *IMC'06*, pages 153–158, 2006.

[8] Z. Mao, D. Johnson, J. Rexford, J. Wang, and R. H. Katz. Scalable and accurate identification of AS-level forwarding paths. In *INFOCOM*, 2004.

[9] Z. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *SIGCOMM*, 2003.

[10] N. Spring and R. Mahajan and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *SIGCOMM*, 2002.

[11] Z. M. Mao, L. Qiu, J. Wang, and Y. Zhang. On AS-level path inference. In *SIGMETRICS'05*, pages 339–349, 2005.

[12] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *ACM SIGCOMM Internet Measurement Workshop 2001*. ACM, 2001.

[13] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie. Sampling biases in ip topology measurements.

[14] CAIDA. Research Project. http://www.caida.org.

[15] University of Oregon RouteViews Project (http://www.routeviews.org/).

[16] C. M. T. P. Team. Caida skitter topology traces (collection).

[17] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of Internet topology using k-shell decomposition. *Proceedings of the National Academy of Sciences USA (PNAS)*, 104(27):11150–11154, July 3 2007.

[18] Y. Shavitt, E. Shir, and U. Weinsberg Near-Deterministic Inference of AS Relationships In arXiv:0711.4562v1, Nov. 2007.