What is Happening from Behind? — Making the Impact of Internet Topology Visible *

Matthias Wählisch[1,2], Thomas C. Schmidt[1] and Waldemar Spät[1]
{waehlisch, t.schmidt}@ieee.org, waldemar.spaet@haw-hamburg.de
[1]HAW Hamburg, Department Informatik, Berliner Tor 7, 20099 Hamburg, Germany
[2]link-lab, Hönower Str. 35, 10318 Berlin, Germany

Abstract

The Internet topology has evolved over the past decades in an evolutionary process and continues to grow. Recently, it has attracted much attention from the networking and physics communities, as it forms a unique operational instance of a planetary-scale network environment. Several measurement projects observing the Internet have been undertaken over the past years, out of which Skitter, its successor Archipelago (Ark) and Dimes have established as continuous recordings of the vivid process of network formation.

In this paper we compare Internet measurement data obtained from Skitter, Ark and Dimes by analyzing the Internet node degree distributions and correlations at IP node and router level. This comparative analysis was enabled by a data conversion and processing tool-chain implemented as an extension to the BRITE topology generator which we introduce, as well. Our results show significant differences in higher nodal degrees. Correlation analysis indicates that DIMES scans discover Internet links to a fairly uniform degree, while parts remain invisible within Skitter and Ark data. Mid-range, oscillating spatial autocorrelations are discovered as a signature of memory effects in Internet topology.

Further on we analyze implications of the Internet structure as attained in both, its core and edge vicinities. Mobile multicast routing performance is quantized by the number of states minimally required for servicing listener or sender mobility. Results show a surprisingly low mobility overhead as compared to general multicast forwarding state management. As continuous mobility handovers necessarily occur between access routers located in geographic vicinity, we finally investigate the hypothesis that geographically adjacent edge networks attain a reduced network distances as compared to arbitrary Internet nodes. The evaluation of edge distance distributions in different regions for IP ranges, clustered according to their geographic location, reveals a stable correlation of geographic and network proximity at Internet edges.

Keywords: Internet topology, measurement, correlation analysis, topology data management, mobility handover performance, inter-domain multicast routing

1 Introduction

The Internet forms a large, decentrally evolving network instance and as such has been subject to many studies. Internet topology models have been developed and suggest that viewing the Internet at different degrees of resolution will lead to scale-free characteristic properties. Empirical measurements of the Internet have thus gained importance for verifying these structural models.

The use of real-world Internet topology data, though, is not limited to theoretical considerations, but is likewise important for protocol design, performance analysis and verifications through realistic simulations. Many protocols such as mobility management or multicast routing on distribution trees,

*This work is supported by the German Bundesministerium für Bildung und Forschung within the project *Moviecast* — http://moviecast.realmv6.org.

overlay routing or hybrid approaches inherit significant performance impact from the underlying Internet topology. Their analysis requires large topology datasets at the IP router level to extract protocol behavior and realistic scaling properties from simulation. To achieve this goal, comprehensive data from continuous Internet monitoring is needed, as well as a feasible tool-chain for processing these voluminous amounts of raw measurements and feeding them into common network simulators.

Several projects have measured the Internet during the past decade, but many of them provided only snapshot views. In contrast, the Skitter work of CAIDA (Ski, 2007) has provided continuous Internet monitoring since 1998 and collected a valuable amount of data. Skitter issues network scans in a traceroute-like fashion from 26 globally distributed monitor points, targeting at a carefully maintained destination list of approximately 971.000 nodes. Recently, the Skitter measurements have been replaced by a successor, the Archipelago project (Ark, 2008). Based on an improved traceroute probing facility 'Scamper', Ark follows a team-probing approach, in which destinations are randomly contacted by one out of a varying number of monitor points within an 48 hours probing cycle. As compared to Skitter, Ark attempts to measure a larger number of nodes, one for each of about 7 million /24 prefixes, but reduces probing to one monitor point per cycle. In general, CAIDA's fairly static approach is tied to a largely controlled experimental environment and allows for liable analysis of the temporal Internet development. The relatively new DIMES initiative (Shavitt and Shir, 2005) proceeds along the line of voluntary monitors. A randomly distributed number of currently more than 15.000 agents perform traceroute scans according to a dynamic destination list. Like in the early SETI@Home, peers submit their scans to a central site, where they are post-processed and aggregated to form a daily snapshot view of the Internet.

In this work we start from the latter monitoring projects and introduce a tool-chain based on extensions we implemented for the BRITE (Medina et al., 2001) topology generator. With its help, we compare structural properties of the datasets to gain insight into their specific characteristics in section 2. Concentrating on node degrees, we find clear indications from a correlation analysis that the overall discovery of Internet links remains incomplete within the Skitter and Ark data, whereas corresponding results for DIMES topology measurements show a fairly exhaustive connectivity. Furthermore, we evaluate and discuss the spatial autocorrelation function as a distinct signature of the graphs.

Grounded on these quantitatively assessed data, we systematically derive implications for inherent measurements of the routing complexity in multicast mobility in section 3, which show a surprisingly low mobility overhead as compared to general multicast forwarding state management. Subsequently we evaluate and analyze edge distance distributions in geographical regions for clustered IP ranges such as a city to quantify the handover performance of Internet mobility management in section 4. As continuous mobility handovers necessarily occur between access routers located in geographic vicinity, the hypothesis is investigated that geographically adjacent edge networks attain a reduced network distances as compared to arbitrary Internet nodes. Conclusions and an outlook finalize the paper in section 5.

2 Topology Data Sets: Comparing Skitter, Ark and DIMES using BRITE

2.1 BRITE Extension & Topology Data Sets

For substantial topology-based analysis and simulation purposes, it is desirable to have at hand a comprehensive tool-chain. In this task, data processing and its structural analysis remain tightly coupled. Converting functions and filter schemas thus should be integrated into an existing topology generator. Such a flexible and extendable generator is BRITE (Medina et al., 2001), which is based on a modular concept with a wide range of topology models and exchange formats. Unfortunately, BRITE is no longer supported and some of the proposed interoperabilities, e.g., with Skitter are not fully functional. In the remaining section we will describe our extensions to BRITE.

We upgraded BRITE to support the im- and export of an IP- or AS-level graph created by Skitter or

DIMES. The raw IP-level Skitter measurements are provided in a specific DB format by CAIDA, which needs to be converted into readable ASCII using CAIDA tools. The BRITE Extension decomposes the traceroute paths in separate, bi-directional edges. The import routine creates a complete graph over all selected Skitter files taken from different monitor points. Thereby all edges are extracted from all input files to shape the final graph. Consequently, including more monitor point data will provide a more complete view of interconnects, which can change the graph characteristic significantly. In the case where two or more measurements for an edge connecting the same vertices are obtained from different monitor files, the mean value of the delay will be assigned as link weight. Whenever the delay between nodes cannot be calculated based on the RTT measurements, a configurable default value will be assumed.

In contrast to Skitter, DIMES aggregates all monitor information into one node and one edge file respectively. Like in the Skitter case, multiply measured edges may be part of the data, which will be resolved analogously. It is important to note, that DIMES data must be pre-processed to revise incompletely written node identifiers. Unresolved router addresses are composed of (last known) source and (first known) destination IP address and a pseudo hop number. If unresolved, this data will cause one IP level node to be mapped onto different, artificial addresses, leading to a distorted hop identification within the graph.

The immense amount of data obtained from the scanning projects cannot be used in full for topology-based simulations, e.g., in a network simulator platform. To generate computationally less expensive reduced topologies of preserved properties, the BRITE Extension also provides two filter schemas:

The Map Sampling algorithm (Magoni and Pansiot, 2002) creates a subgraph with a predefined number of nodes and a mean degree reflecting the input topology.

The Radial Neighborhood View extracts a subset of the original topology centered within a predefined hop distance around a preselected node. Let $G = (V, E)$ be the original topology and $v \in V$ the chosen center node, then this view selects the subgraph $G' = (V', E')$ with $V' = \{v' \in V \mid dist(v, v') \leq \text{maxHop}\}$.

Topology input data are taken as sanitized DIMES data and Skitter IP level measurements aggregated over all monitor points. Our graph calculations are based on the freely available BRITE analysis routines contributed by Mathias Golombek. We extended this analysis tool for calculating the Joint Degree Distribution.[1]

2.2 IP Alias Resolution

Trace collections as obtained from Skitter and Ark are likely to include different IP addresses that belong to the same router. Such aliases need to be resolved prior to constructing a router-level topology. There are several heuristic method for alias discovery, cf. (Donnet and Friedman, 2007) for a comprehensiv overview.

Aliases are not immediately visible from traceroute paths, but can be identified using analytic and active methods. In our analysis we employed a combination of a graph-based method to identify alias candidates analytically, followed by an active probing with the CAIDA tool 'iffinder'[2], which evaluates ICMP error messages for packets sent with IP RECORD ROUTE option.

Dimes data are already aggregated in node and edge files and thus withstand a trace-based alias resolution. However, the project provides router files that identify different IP interfaces belonging to the same router.

[1] All BRITE extensions and the DIMES pre-process script can be downloaded from www.realmv6.org/brite-extension.html.

[2] see http://www.caida.org/tools/measurement/iffinder/

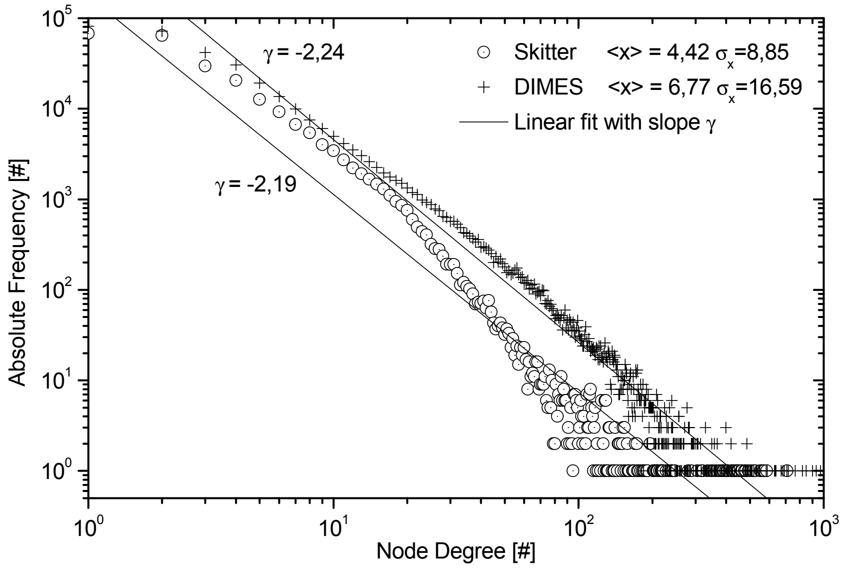

(a) Skitter (raw) versus DIMES

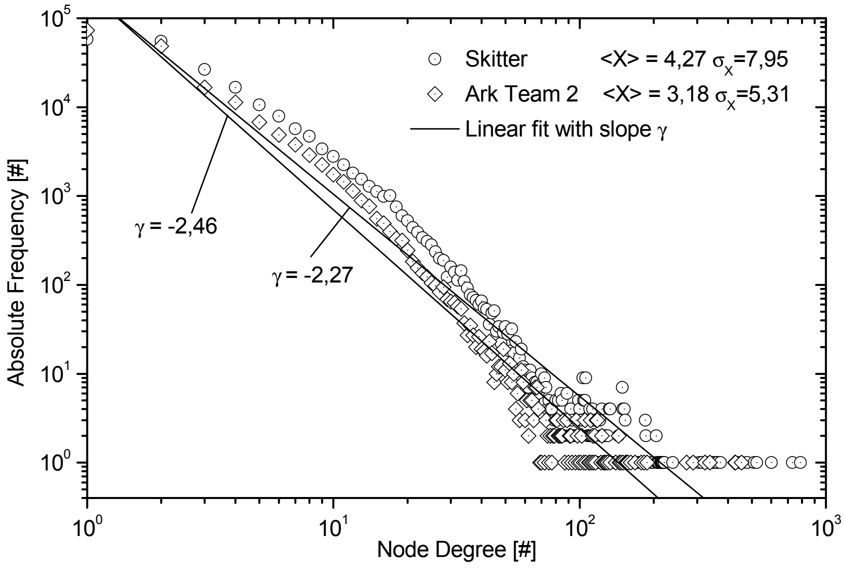

(b) Skitter and Ark with *iffinder* IP Alias resolution

Figure 1: Node Degree Distribution x with mean $<x>$ and dispersion σ_x

2.3 Node Degree Analysis & Results

In this section we apply the obtained tool-chain to a comparative analysis of Skitter and DIMES measurement data. We focus on the basic property of node degree distributions and correlations within the obtained network graphs.

The *degree distributions* for single nodes are displayed in figure 1 on a log-log scale along with a linear best fit of a corresponding power law. We compare raw Skitter data with the DIMES network graphs and apply *iffinder*-based IP Alias resolution to Skitter and Ark measurements. While the resulting power law exponents $-\gamma = 2.19, 2.24$ and 2.27 for Skitter and DIMES remain compatible and close to the original values of Faloutsos et al. (Faloutsos et al., 1999), the corresponding Ark value deviates. In general, the statistical accuracies evidently differ, increasing from $0,947$ for Skitter to $0,970$ for DIMES. The datasets mainly deviate for higher nodal degrees, with Skitter and Ark falling relatively shorter for degrees above 200.

A closer insight into the interconnection properties of the graphs is given by the *joint degree distribution*. This correlation law defines the probability that a randomly selected edge connects nodes with degree k_1 and k_2. Let $m(k_1, k_2)$ denote the number out of M total edges directed from k_1 to k_2 degree nodes, then the correctly normalized joint degree distribution is calculated as

$$P(k_1, k_2) = m(k_1, k_2)/M.$$

It does not only describe the one hop neighborhood structure of an average k-degree node, but can also be used to derive other well-known measures (Newman, 2002; Mahadevan et al., 2006). Note that the single node degree distribution $p(k)$ does not directly follow from integration, but requires a bias correction factor, i.e., $p(k) \propto \sum_j P(k, j)/k$.

Figure 2 displays the corresponding joint degree distributions for Skitter, DIMES and Ark as linear-log contours. For our bi-directionally linked graphs, the probability distribution is symmetric and admits a steep, but smooth decay from low to high node degree correlations. The latter is significantly sharpened for the Skitter and Ark datasets, representing a reduced “visibility” of links between nodes of degrees above 25. Concomitantly, Skitter and Ark distributions admit large, discontinuous wholes in high degree regions. As both measurements where taken at the same time, and as we can expect distributions for the statistical size of the Internet to smoothly fill the event space, these results clearly indicate an incomplete picture or “limited” horizon of the topology discovered by the small number of Skitter monitor points. Ark team probing, which reduces standard tracing of a destination to only one monitor, enhances these effects.

The *assortativity coefficient* r, given in captions of fig. 2, is the correlation coefficient of the neighboring degrees and serves as a simple measure of first order dependence for the network graph. For AS-level topologies, it has been consistently identified as of significant negative values (Newman, 2002; Mahadevan et al., 2006), indicating the disassortative tendency of AS’ to connect at dissimilar degrees. By contrast, the corresponding node-level values for the Skitter and DIMES topologies show an assortative preference of IP routers to interconnect at compatible degrees, while slightly negative values are attained subsequent to alias resolution in Skitter and Ark. However, the correlation values of Skitter and Ark are almost one order of magnitude lower than for DIMES, pointing to further differences in the internal correlations of the graphs obtained from the two datasets.

To gain a deeper understanding of the mesh properties of the graphs, we examine long-range interdependencies as described by the spatial autocorrelation function of node degrees. Let $x(v)$ denote the degree for a node $v \in V$ with average $< x >$ and dispersion σ_x, where averages are performed over the complete topology $G = (V, E)$, and $X(v) = (x(v) - < x >)/\sigma_x$ its centered and normalized transform. Further for a given node v_i define $v_{i,\delta} \in \{v' \in V | \, dist(v_i, v') = \delta\}$ to be a node of hop distance δ, then the spatial autocorrelation function reads

$$\rho(\delta) = < X(v_i) \cdot X(v_{i,\delta}) > .$$

The complex measure $\rho(\delta)$ quantifies the similarity of the node degree distribution with itself at a position shifted by δ hops. Obviously $-1 \leq \rho \leq 1, \;\rho(0) = 1$ and the decay of ρ characterizes the

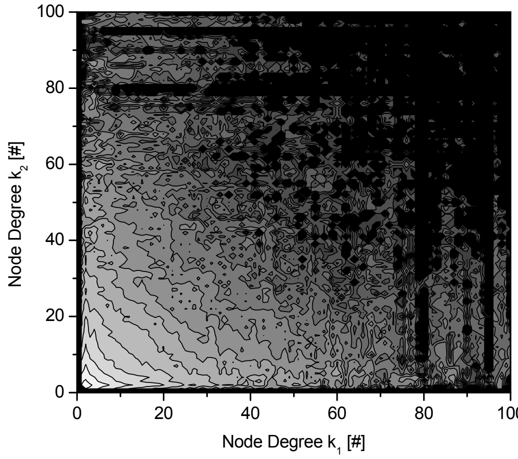

(a) Skitter, Assortativity Coefficient $r = 0,011$

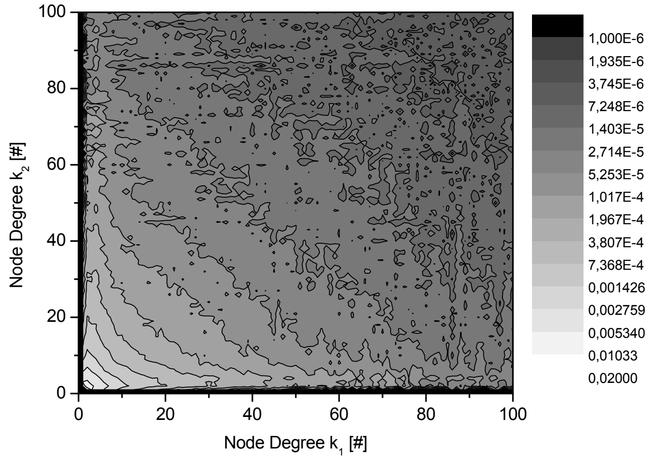

(b) DIMES, Assortativity Coefficient $r = 0,091$

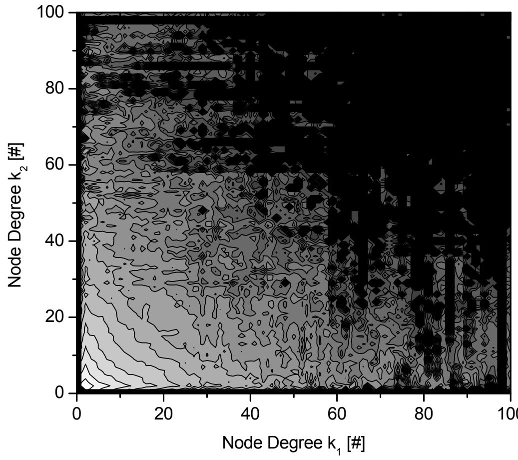

(c) Skitter with *iffinder*, Assortativity Coefficient $r = -0.0094$

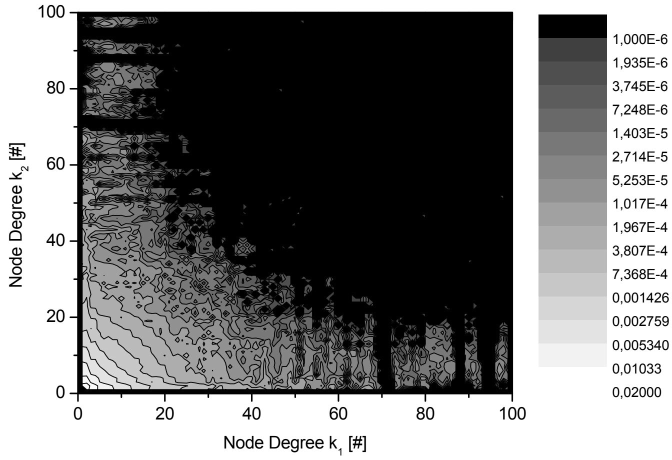

(d) Ark with *iffinder*, Assortativity Coefficient $r = -0,0028$

Figure 2: Linear-Log Contours of the Joint Degree Distributions $P(k_1, k_2)$, for $1 \leq k_i \leq 100$.

memory range of the stationary process X: For a random graph it admits zero values for $\delta > 0$, its exponential decay indicates independent increments, i.e., the degree of a node reached in δ hops changes as compared to its neighbors independent of the hop distance δ. Distributed memory effects are present for polynomial decays, with long-range dependencies dominating for $\sum_\delta |\rho(\delta)| = \infty$.

The corresponding results for Skitter and DIMES data are displayed in figure 3 along with least square fitted analytic curves. This long-range correlation picture reveals further data discrepancies. While the Skitter autocorrelation function quickly decays, values for DIMES decent in polynomial order from positive nearby-neighbor degree correlation and oscillate to a negative regime for hop distances from 5 to 7. These observations appear consistent with anticorrelations previously found at AS-level transitions in (Newman, 2002; Mahadevan et al., 2006). As expected, autocorrelation converges to zero for large distances in DIMES, whereas Skitter values saturate at a slightly negative level. The latter may be contributed to measurement artifacts.

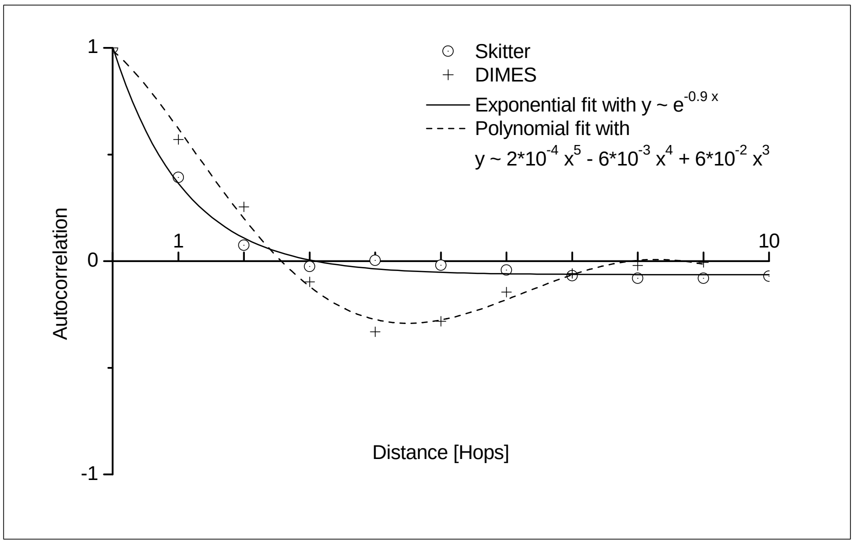

Figure 3: Spatial Autocorrelation Functions

3 Exploring the Routing Complexity of Mobile Multicast

3.1 Problem Analysis

Unlike point-to-point mobility and despite of ten years of active research, mobile multicast protocol development is still in an early, premature state (Schmidt et al., 2008). The complexity of multicast forwarding state management, as already present in static environments, may be recognized as one of the major reasons for hesitant engagement at the IETF. To acchieve optimal routing, any client subscribed to a group while in motion, requires delivery branches to pursue its new location. Any mobile source requests the entire delivery tree to adapt to its changing positions. However, multicast distribution trees arising from handover scenarios are not independent, but highly correlated. It is the aim of this paper to give realistic quantitative estimates on handover-initiated state establishment and thereby to provide a minimal bound on the complexity of mobile multicast routing. Our analysis concentrates on source specific shortest path trees (SPT) and is therefore valid for shared and source specific tree protocols at the receiver side, but restricted to SPT based routing for senders.

The effect of source or receiver movement on the stability of shortest path trees is directly addressed by constructing and comparing multicast distribution trees. As key characteristics of multicast shortest path trees only make an impact in large networks, and as topological setup fixes a dominant part of the degrees of freedom in routing simulations, realistic Internet topology data is needed for our analysis, which we obtain from the data and tool-chain presented in the previous section.

The multicast tree analysis has been realized based on the network simulator OMNeT++ 3.3 (Varga et al., 2007). In detail, we uniformly sample receivers and sources as attached to, and designated multicast routers from the edge nodes of the given topology data sets. Edge nodes are identified as routers of degree one and represent transition points to 'customer' networks within the Internet core systems. A source resp. *one* receiver node is selected as mobile and moved from previous (pDR) to next designated router (nDR). The routing distance from pDR to nDR forms a basic measure of mobility impact and is chosen at varying but predefined hop count. The next distribution tree is then spanned from the nDR to all receivers in the mobile source scenario and with altered branch on the path from source to the mobile in the receiver scenario. These trees are then analyzed w.r.t. changes in forwarding interfaces.

3.2 Semi-Empirical Results

Results for the relative change of distribution trees as a function of receiver multiplicity are shown in figure 4 for a step size of 5 (4(a), 4(b)) and 10 hops (4(d), 4(e)), as well as a function of DR-distances (4(c)). It is noteworthy that even in large networks and for moderate receiver numbers more than 80% of multicast tree routes remain fixed under a handover.

An increasing number of receivers within the network will broaden its coverage with distribution states and lead to lower state changes under mobility (cf. 4(a) - 4(e)). Conversely, the decreasing stability of multicast SPTs for increasing mobility impact is clearly visible from figure 4(c). When comparing mobility scenarios, it is noticeable that network size significantly influences multicast tree stability only for moving receivers, due to the relevance of receiver-to-overall-node ratio. Source movement remains fairly unaffected, which suggests a scaling law argument, i.e., a self-similar nature of multicast (sub-)trees added at distance in larger networks.

The coincidence rate for receivers remains below the results for mobile sources. Even though a moving source changes the tree root, its leaves represented by fixed receivers remain unchanged and advance the correlation of mobility related distribution trees. This scenario does not apply for mobile receivers and the results suggest that the intersection between the old and the new tree is located close to the mobile. In the overall it can be concluded that the stability of multicast distribution trees under mobility is surprisingly large, which raises hope for an efficient solution in the near future.

4 Handover Performance of Internet Mobility Management

Seamless support for Voice over IP (VoIP) and related real–time communication must be considered critical for deployment success into the mobile world. Therefore significant effort is continuously taken in the IETF to develop and improve protocols for seamless mobility handovers, FMIPv6 (Koodli, 2008) and HMIPv6 (Soliman et al., 2005) being the most prominent examples. Mobile IPv6 (Johnson et al., 2004) inherits a strong topology dependence through its binding update procedures with the Home Agent (HA) and the Correspondent Node (CN). Handover acceleration schemes attempt to overcome this obstacle by relocating immediate transfer negotiations to the vicinity of the mobile node, i.e., to access networks at the Internet edges. In previous analysis (Schmidt and Wählisch, 2005) it could be shown that the handover performance actually attained largely depends on the relative network topology of access components, when measured in an appropriate delay metric such as round trip time. Access router distance can be considered as the characteristic complexity parameter in fast or hierarchical mobile IPv6.

4.1 Regional Edge Distance Distributions

In this section we empirically analyze the regional edge distance distributions of exemplary areas in the current Internet. The objective of this work is to inquire on the temporal handover distribution, which is an immediate function of access router distances. As the measure of locality admits enhanced sensitivity on local edge network structures, we performed own network scans to allow for more detailed investigation on regional topologies. Clusters of IP ranges from geographic regions such as cities are pre-selected in order to account for locality. To choose for a reliable source of geographic information, we evaluated eight different mapping resources in a first step by selecting a set of 30 distributed, geographically known IP ranges. The commercial product GeoIP (MaxMind LLC, 2006) thereby was the only resource to admit negligible errors. In a second, automated testing we compared data of larger samples with *whois* queries and found a coincidence rate of about 80 %. This result we considered reasonable, as *whois* data commonly provide administrative addresses possibly distinct of physical router locations.

Results are presented for the cities of San Francisco, USA, Berlin and Hamburg, Germany, and Shanghai, China, which were exemplarily selected as geographic target regions. Scanning has been performed from origins at the locations of Berlin, Hamburg and Shanghai, at the 67th IETF meeting

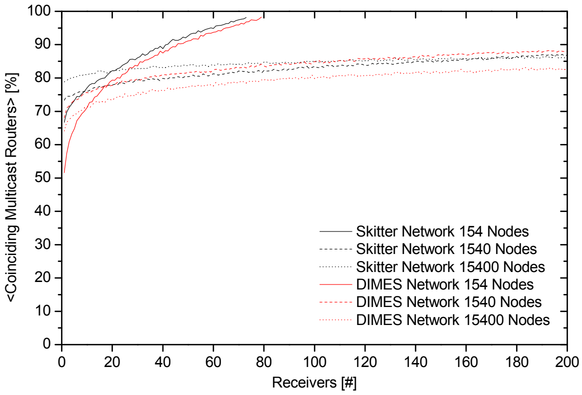

(a) Mobile receiver

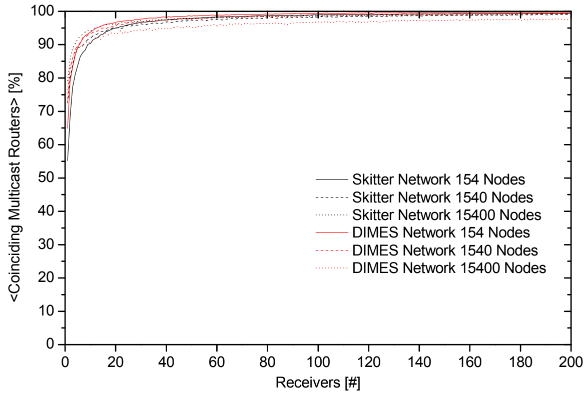

(b) Mobile sender

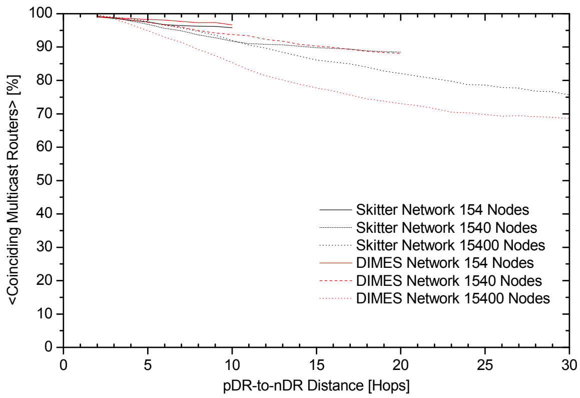

(c) Mobile sender

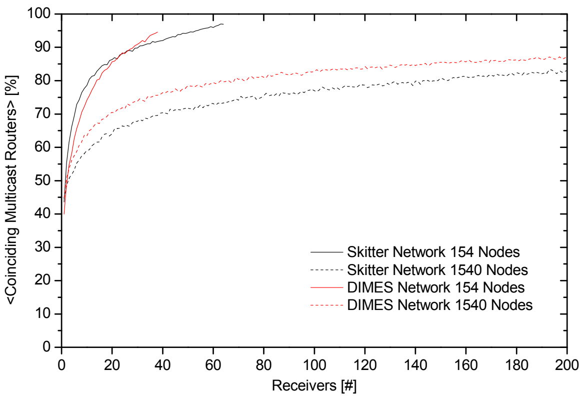

(d) Mobile receiver

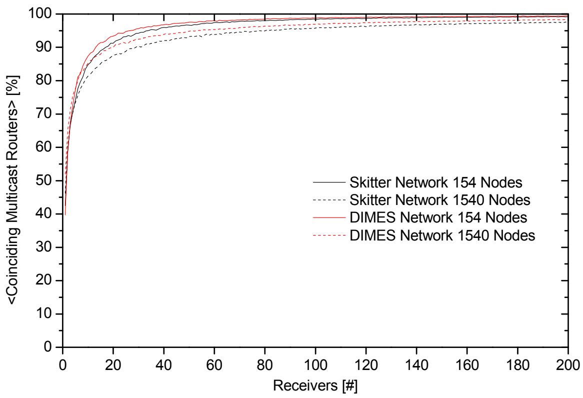

(e) Mobile sender

Figure 4: Relative router coincidence for subsequent multicast trees at pDR-nDR distances 5 & 10

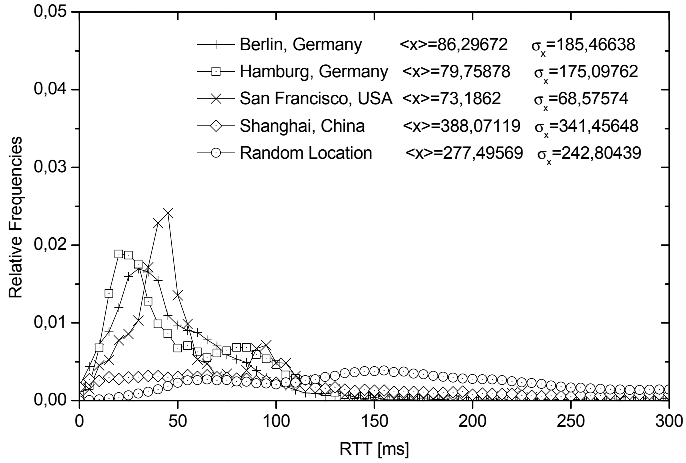

(a) Scan Data

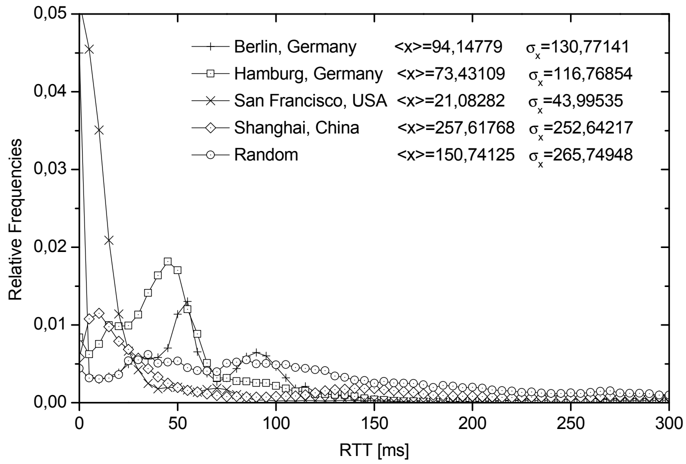

(b) CAIDA Data

Figure 5: Round Trip Time Distributions at Network Edges

in San Diego and with the help of various public traceroute facilities. Since the number of available IP ranges vary from Shanghai (763) up to San Francisco (8476), subsets of equal sizes are selected randomly for each city. Statistical convergence with respect to sample size, but also for different dates and day times were compared, and a fair stability of the distributions could be observed for sample subsets of 500 IP ranges.

We compare our results with distributions derived from CAIDA data. DIMES data do not include RTT values and are thus unavailable. Host clusters for selected cities are taken from the CAIDA destination list according to the GeoIP database. Skitter trace paths are minimized with respect to all available 18 monitor points, which are located more densely at the US West Cost and sparsely in Europe and Asia.

Resulting distributions for round trip times are displayed in figure 5. An additional curve derived for randomly located nodes is added to distinguish locality correlations.

Clearly our round trip time distributions vary significantly, while Skitter data evaluate to fairly similar distributions. Pronounced peaks at close distances can be observed for the areas of Berlin and Hamburg, when monitored from the close vicinity. The effect on RTT results of source positioning for the scanning is shown for San Francisco data in figure 6. RTT characteristics, though, appear heavier tailed, which supposedly is due to a sporadic occurrence of slow transition links. Tardy transitions are of lesser effect in the San Francisco region, for which again CAIDA measurements expose a distribution of higher significance.

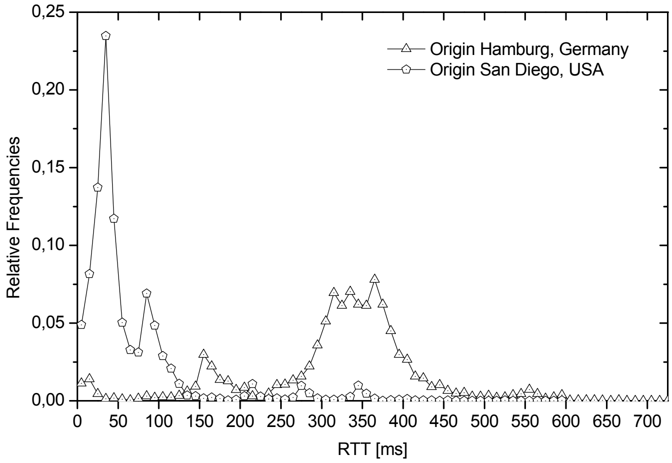

Figure 6: RTT Distributions Evaluated from Different Origins for San Francisco Nodes

In contrast, Shanghai data attain merely indifferent RTT distributions, which are even less pronounced than random samples. Non-negligible weights are situated beyond the displayed interval, as can be read from mean and standard deviation values. This may be explained from a wide variety of slow transit links present in the Chinese core networks.[3] CAIDA skitter data seem to qualitatively reflect these RTT law diversities. San Francisco values are very pronounced, whereas Hamburg and Berlin data show an intermediate characteristic. It should be noted, though, that the reverse of the proximity observations from (Subramanian et al., 2002) does not seem to hold: RTT distributions admit wide tails, whence even in close router distances enhanced mutual delays may be expected.

In total the results seem to indicate that inter–edge routing within a geographic region is frequently performed via local transits and peering, which produce network proximity in 'the neighborhood', but remain invisible for a distant monitor. From an applied perspective, these results support the hypothesis of section 2 that Internet link correlations only become visible when observed from an appropriately close perspective.

4.2 Internet Handover Performance

The results obtained so far may serve as an empirical fundament for realistic handover performance estimates of the network. A mobile node moving from one access network to another in geographical neighborhood does imposes traffic redirection, minimally from its previous to its new attachment. These operations cause delay and routing costs, which for the case of FMIPv6 (Koodli, 2008) are given by the unicast path from previous to next access router, and higher, otherwise.

Based on the results derived in (Schmidt and Wählisch, 2005), we now can immediately calculate expected values of characteristic handover measures. For packets sent at a constant bit rate of one per 10 ms, the conditional expectation of packets lost or buffered for given inter–access–router delay was derived for predictive and reactive handover procedures (cf. figure 6 of (Schmidt and Wählisch, 2005)). Combining these previous results with those shown in figure 5, we arrive at expected periods for packet loss as functions of handover anticipation times. Results for the different regions as presented in figure 7 jointly show a pronounced uniform minimum at handover anticipation of 25 ms for the cities of San Francisco, Hamburg and Berlin, while significant optimal values remain absent for Shanghai and random data. These results reflect the degree of locality in regional delay distributions.

[3] Another possible explanation could lie in a reduced accuracy of MaxMind GeoIP data for the Chinese region.

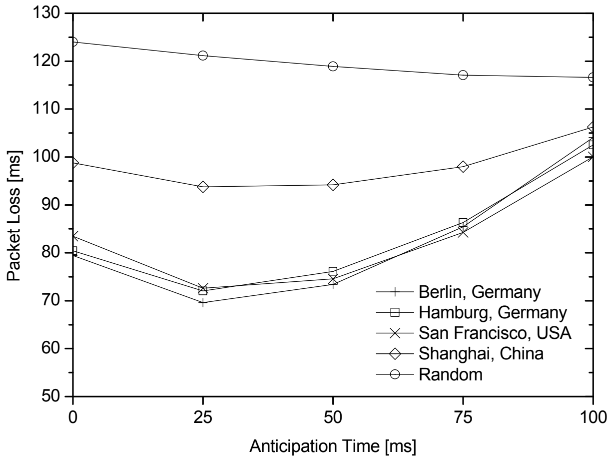

Figure 7: Expected Packet Loss in Predictive Handovers as a Function of Anticipation Time (0 = Reactive Handover)

5 Conclusions & Outlook

We introduced a tool-chain for data processing of Internet measurements and used it for a comparative analysis of CAIDA Skitter, Ark and DIMES data on the IP node level. Topological characteristics of the two datasets could be derived from a node degree correlation analysis, indicating a fairly uniform discovery of Internet links by DIMES, but a 'visible' miss by Skitter and Ark. Investigations of long-range autocorrelations indicate that the Internet topology shows mid-range memory effects, with preferred clustering of similar degree nodes in the short range, but dissimilar ones on an inter-AS level scale.

These data sets and tools where applied to analyze aspects of mobility management at the routing layer. At first we quantified the complexity inherent to multicast tree adaptation under mobility. Topology-based simulations for CAIDA and DIMES data sets discovered a remarkably low requirement of multicast state change in moderate mobility regimes.

Inspired by mobile IP handover performance measures, we investigated routing distances in geographically bound clusters of the Internet. Therein users are expected to move around freely while continuously 'talking' IP in the near future. Traceroute probes have been used to derive delay distributions at Internet edges in San Francisco, Berlin, Hamburg and Shanghai. Comparison has been drawn to CAIDA measurements only, since DIMES link data lack delay values. Our results seem to indicate a clear signatures of locality in the distance metric, which cannot be segregated from CAIDA measurements due to sparsely scattered monitor points. The application of these results to calculating packet loss after mobility handovers indicates that characteristic proximity measures in the Internet may give rise to fairly stable anticipation timers.

References

Ark (2008), 'Archipelago Project at CAIDA', http://www.caida.org/projects/ark/.

Donnet, B. and Friedman, T. (2007), 'Internet Topology Discovery: A Survey', *Communications Surveys & Tutorials, IEEE* **9**(4), 56–69.
URL: *http://dx.doi.org/10.1109/COMST.2007.4444750*

Faloutsos, M., Faloutsos, P. and Faloutsos, C. (1999), On Power-Law Relationships of the Internet Topology, *in* 'SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication', ACM Press, New York, NY, USA, pp. 251–262.

Johnson, D. B., Perkins, C. and Arkko, J. (2004), Mobility Support in IPv6, RFC 3775, IETF.

Koodli, R. (2008), Mobile IPv6 Fast Handovers, RFC 5268, IETF.

Magoni, D. and Pansiot, J.-J. (2002), Internet Topology Modeler Based on Map Sampling, *in* 'Proceedings of the 7th IEEE Symposium on Computers and Communications', IEEE Computer Society, Taormina, Italy, pp. 1021–1027.

Mahadevan, P., Krioukov, D., Fomenkov, M., Huffaker, B., Dimitropoulos, X., kc claffy and Vahdat, A. (2006), 'The Internet AS-Level Topology: Three Data Sources and One Definitive Metric', *ACM SIGCOMM Comput. Commun. Rev.* **36**(1), 17–26.

MaxMind LLC (2006), 'GeoIP', http://www.maxmind.com.

Medina, A., Lakhina, A., Matta, I. and Byers, J. (2001), BRITE: An Approach to Universal Topology Generation, *in* 'Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems', IEEE Computer Society, Washington, DC, USA, p. 346.

Newman, M. E. J. (2002), 'Assortative Mixing in Networks', *Physical Review Letters* **89**(20), 208701.

Schmidt, T. C. and Wählisch, M. (2005), 'Predictive versus Reactive – Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility', *Telecommunication Systems* **30**(1–3), 123–142.

Schmidt, T. C., Wählisch, M. and Fairhurst, G. (2008), Multicast Mobility in MIPv6: Problem Statement and Brief Survey, IRTF Internet Draft – work in progress 04, MobOpts.
URL: *http://tools.ietf.org/html/draft-irtf-mobopts-mmcastv6-ps*

Shavitt, Y. and Shir, E. (2005), 'DIMES: Let the Internet Measure Itself', *ACM SIGCOMM Computer Communication Review* **35**(5), 71–74.

Ski (2007), 'Skitter Project at CAIDA', http://www.caida.org/tools/measurement/skitter.

Soliman, H., Castelluccia, C., Malki, K. and Bellier, L. (2005), Hierarchical Mobile IPv6 mobility management (HMIPv6), RFC 4140, IETF.

Subramanian, L., Padmanabhan, V. and Katz, R. (2002), Geographic Properties of Internet Routing, *in* 'Proceedings of the 2002 USENIX Annual Technical Conference', USENIX Association, Berkeley, CA, USA, pp. 243–259.

Varga, A. et al. (2007), 'The OMNeT++ discrete event simulation system', http://www.omnetpp.org.