

# Geographic Locality of IP Prefixes

Michael J. Freedman  
New York University  
mfreed@cs.nyu.edu

Mythili Vutukuru, Nick Feamster, Hari Balakrishnan  
Massachusetts Institute of Technology  
{mythili,feamster,hari}@csail.mit.edu

## ABSTRACT

Information about the geographic locality of IP prefixes can be useful for understanding the issues related to IP address allocation, aggregation, and BGP routing table growth. In this paper, we use traceroute data and geographic mappings of IP addresses to study the geographic properties of IP prefixes and their implications on Internet routing. We find that (1) IP prefixes may be too coarse-grained for expressing routing policies, (2) address allocation policies and the granularity of routing contribute significantly to routing table size, and (3) not considering the geographic diversity of contiguous prefixes may result in overestimating the opportunities for aggregation in the BGP routing table.

## 1. Introduction

Today’s Internet routing infrastructure achieves scalability by expressing reachability for large groups of IP addresses using a single IP prefix in a route advertisement. Today’s largest Internet routing tables provide reachability to hundreds of millions of end hosts with nearly 200,000 routes [5]. IP addresses that are nearby in IP space may be geographically or topologically diverse, and vice versa. This paper *quantifies* this lack of correspondence. Information about the geographic location of hosts within IP prefixes can also help us better understand many issues related to IP address aggregation and allocation and their effect on BGP routing table growth.

Our study uses extensive traceroutes and leverages IP-to-geographic mapping techniques to examine the geographic properties of multiple destinations *within a single prefix*. Our dataset includes traceroutes to at least 4 IP addresses within each prefix of the global routing table, as well as traceroutes to 1.6 million unique Web clients and servers that exchanged content over CoralCDN, a popular peer-to-peer content distribution network [3].

Towards this goal of understanding the geographic properties of IP prefixes, this paper makes three findings. First, an IP prefix may express only very coarse geographic information about the destinations (and networks) that it comprises. This property of the geographic diversity of hosts within a prefix is important for techniques that assume that hosts within an IP prefix are topologically close.

As expected, we find that “shorter” IP prefixes, which represent a larger portion of the IP address space, tend to comprise destinations in a large number of geographic locations, spread over long distances. For example, more than half the prefixes with mask lengths between 8 and 15 span a distance of more than 100 miles. More surprisingly, we find that “longer” prefixes, albeit a small fraction of them, can be quite geographically diverse: about 1.4% of the prefixes with mask lengths between 24 and 31 span a distance of more than 100 miles, and some /24 prefixes span distances of more than 10,000 miles!

Second, autonomous systems (ASes) commonly advertise multiple discontinuous IP prefixes for networks in the *same* geographic location. In this case, the Internet routing table must carry multiple routes for a group of destinations in a single geographic location and a single AS, because the addresses cannot be expressed as a single IP prefix. This finding suggests that an Internet routing infrastructure whose routing granularity more closely reflects geography could significantly reduce the size of the global routing tables. Additionally, fragmented address allocation explains 65% of the cases where a single AS was advertising discontinuous prefixes from the same location, which suggests that IP address renumbering could significantly reduce the size of the BGP routing table.

Finally, ASes sometimes announce contiguous prefixes from *different* geographic locations. Ongoing studies, such as the CIDR Report [2], presume that all contiguous prefixes originated by an AS should be aggregated into a single IP prefix. However, these studies do not consider whether these prefixes actually represent geographically diverse networks that are intentionally represented as separate routes. By ignoring location information, the CIDR Report may overestimate the opportunities for aggregation by a factor of three.

## 2. Related Work

Padmanabhan *et al.* [9] develop a set of techniques to map IP addresses to geographic locations. One of their techniques “clusters” IP addresses at the granularity of an IP prefix to map them to a location. The authors observe that the accuracy of their method in mapping an IP address is related to the geographic spread of the hosts within the

prefix containing that IP address. Our work aims to gain a deeper understanding of geographic diversity of the hosts within a single IP prefix.

The geographic locality of IP prefixes is significant for systems like Network Aware Clustering (NAC) [6], which group hosts that belong to the same prefix of the BGP routing tables into clusters, which are used in applications like content distribution and proxy positioning. These clustering schemes rely on the assumption that hosts within a prefix are likely to be topologically close and under the same administrative domain. We investigate the validity of this assumption in Section 4.1.

Earlier work has also studied impact of factors like IPv4 address allocation and aggregation on the growth of the BGP routing table [1, 7]. Bu *et al.* [1] find that address fragmentation (where a set of prefixes originated by an AS *cannot* be summarized by one prefix) is the biggest factor contributing to BGP routing table growth. Our study also reveals many instances where an AS announces discontinuous prefixes, even from the *same* geographic location.

The CIDR Report studies contiguous prefixes announced by the same AS and the missed opportunities for aggregation by ASes [2]. In our study, we find that contiguous prefixes announced by the same AS are sometimes geographically far apart; aggregating such prefixes might conflict with an AS’s traffic engineering or load balancing goals. Thus, the aggregation opportunities suggested by the CIDR Report might not all be feasible.

### 3. Data

This paper uses three datasets generated by traceroute measurements to study the relationship between IP prefixes and locality. We mapped IP addresses to IP prefixes using longest-prefix matching on a BGP table from RouteViews [8] from February 27, 2005. This table had approximately 170,000 IP prefixes.

As shown in Table 1, *Clients* and *Servers* refer to traceroutes taken to Web clients and servers that exchanged content over CoralCDN, a peer-to-peer content distribution network that receives approximately 10 million HTTP requests per day from widely-dispersed clients [3]. The client traces cover a 14-day period starting on February 13, 2005, while the server trace covers a single day (April 26, 2005). Each CoralCDN Web proxy—there are approximately 225 such proxies deployed on PlanetLab [10]—performed a traceroute to every client *destination* IP.

While these CoralCDN datasets provide a workload corresponding to a real user population, we also sought to provide coverage of all IP prefixes from the RouteViews table. For the *Breadth* dataset, we performed traceroutes to 4 uniformly distributed IP addresses per advertised prefix, using 25 PlanetLab hosts as sources. Note that these traceroutes traverse IP addresses from multiple prefixes.

Dataset	Period	Traceroutes	Destinations	IPs	Prefixes
Clients	Feb 13-27, 2005	6,565,844	1,599,228	692,080	45,573
Servers	Apr 26, 2005	71,621	36,387	64,378	9,589
Breadth	Apr 25, 2005	675,797	649,441	246,626	161,974

**Table 1: Traceroute datasets. The last two columns show *reachable* IP addresses and prefixes: routers and destinations from which ICMP replies were received.**

Thus, *Breadth* actually includes many more data points than four per prefix, especially for transit ASes.

Dataset	Mapped	Inherited	Prefixes	ASes	Locations
Clients	313,573	180,487	6,136	1,244	1,363
Servers	22,749	5,032	1,693	541	748
Breadth	176,601	130,621	6,828	1,605	1,206

**Table 2: IP-to-location assignments.**

We use the RouteViews table to map IP addresses to their ASes and DNS naming heuristics to map IPs to locations, as described in Section 3.1. Table 2 characterizes the number of IP addresses *mapped* to an AS number and a location (at the city level). We call this location *inherited* if the destination is not reachable itself (whereupon we assign it to the location of its closest reachable upstream router instead). The *inherited* dataset is a subset of *mapped*, which in turn is a subset of the *destination* IPs in Table 1. Table 2 also shows the total number of unique IP prefixes, ASes, and locations in each dataset.

#### 3.1 Mapping IP addresses to locations

We use `undns` [11] to map IP addresses to locations. `undns` extracts geographic information from a DNS name, which is useful because network operators often use geographically meaningful names for routers. For example, a DNS name of the form *qwest-gw.n54ny.ip.att.net* refers to an AT&T (AS 7018) router peering with Qwest, located at an exchange point on 54th street in New York City. Other studies have also used this approach [9].

Unfortunately, naming heuristics vary between ISPs, and parsing is a manual process. ISPs may name routers by city name or code, airport code, or some 4-to-6 letter abbreviation for city and state. In addition, ISPs incorporate such information in hostnames differently; even a single AS may use multiple heuristics. For example, Verio (AS 2914) names gateways in one manner (*e.g.*, *att-gw.nyc.verio.net*) and customer addresses in another (*e.g.*, *vl-101.a02.nycmny03.us.ce.verio.net*). Router names can also be ambiguous: for example, *nycmng-washng.abilene.ucaid.edu* is located in New York but peers with a router in Washington, D.C. In such special cases, we manually pinged routers from diverse locations to better understand their ISP-specific naming heuristics.

`undns` version 0.1.27a includes manually written hostname parsing rules for 247 ASes, mostly Tier-1 and Tier-2 ISPs in the US and Europe. We added support for 169 additional ASes (including smaller ISPs) and expanded the

tool’s international coverage. The latter is especially important for the *Clients* dataset, which includes significant amounts of traffic from Asia. We spot-checked location estimates after running `undns` for some IP addresses in known locations.

Given a city-level location estimate for a particular IP address, we also assign to it the latitude and longitude coordinates for that city, which allows us to estimate the distance between two IP addresses.

### 3.2 Limitations of mapping technique

Our data has several limitations. First, a reverse DNS mapping from IP address to hostname may not exist; such records existed for only 50%-60% of all *unique* reachable IP addresses. Second, `undns` may not have a parsing rule to map the hostname to a location; our ruleset assigned locations to about one-third of known hostnames. Third, `undns` may return incorrect IP-to-AS number mappings. Finally, some destinations were not reachable via traceroute. We now discuss mitigating factors for the first two limitations and solutions for the latter two.

While we could resolve the hostnames of less than 60% of IPs, we found that internal ISP routers—as opposed to gateway routers or customer addresses—were more commonly missing reverse DNS records. These routers are unlikely to express more geographic diversity than that already captured by gateways and customers, so this limitation should not significantly affect our results.

Even though `undns` assigned locations for only one-third of all *unique* hostnames, two factors reduced the impact of this poorer coverage. First, our ruleset provides very good coverage for real-world traffic patterns, as we supply more detailed rules for popular ASes. In fact, we resolved the location of 90% of probed IPs in *Servers* (*i.e.*, when counting *all* instances, instead of only unique instances, of hostnames). Second, the hostnames that had no locality information were most commonly at the network edges where dynamic addressing is used (*e.g.*, cable modem, DSL, and dialup connections). This may inflate the number of hosts with unassigned locations.

`undns` uses the hostname of an IP address to determine its AS number, which could cause us to mistakenly believe an ISP is announcing a discontinuous prefix. For example, an IP address in AS 6395 (Broadwing Communications) carries the hostname suffix *.northwestern.edu*, even though its corresponding /14 prefix is announced solely by Broadwing, which provides transit service for Northwestern University (AS 103). To solve this problem, we assigned an AS number to an IP address by performing longest-prefix matching against the RouteViews table.

Finally, many destinations were not directly reachable when performing traceroutes: 57% of addresses in *Clients*, 52% in *Servers*, and 76% in *Breadth*. This limitation is generally due to firewalls blocking ICMP packets at large

portions of the networks’ edges. and many destinations in *Breadth* were unused IP addresses. To solve this problem, we assigned an unreachable destination IP address to the location of its last reachable upstream router. Our use of traceroutes enables us both to discover routable IP addresses for firewalled or unused destinations and to determine the upstream addresses for inherited locations.

## 4. Results

We first examine the geographic diversity of individual IP prefixes, paying particular attention to the *maximum* geographic distance between any two pairs of IP addresses within a single prefix. We then study the extent to which a single AS advertises multiple discontinuous prefixes that refer to endpoints at a single location, as well as the causes of these advertisements. Finally, we study the extent to which an AS advertises contiguous prefixes for hosts in diverse geographic locations.


### 4.1 Single prefix with multiple locations

In this section, we study the extent to which a single IP prefix comprises hosts in multiple geographic locations (thus potentially obscuring potentially useful information by over-aggressive aggregation). Figure 1(a) shows the number of distinct geographic locations contained within a single geographic prefix for the *Clients* dataset. As expected, shorter prefixes tend to comprise more geographic locations.

Figure 1(b) shows that, not only do the shorter prefixes span more geographic locations, but these hosts also span a much wider geographic distance: nearly half of the prefixes in the /8-/15 range span a distance of more than 100 miles. Several of the prefixes in this range are either European backbones or broadband access providers in the United States: for example, from the *Clients* dataset, we find that AS 7132 (SBC) advertises a single /16 that contains 64 distinct locations spread across the United States. Transit ASes with smaller address allocations also advertised prefixes containing geographically diverse hosts: *e.g.*, AS 7657 (The Internet Group, a New Zealand ISP), advertised a /24 whose IP addresses span 1,400 miles.

Because ASes (particularly US-based backbone ISPs) often allocate sub-prefixes from a single large IP prefix, we expected that prefixes that are allocated to transit ISPs are more likely to have geographically diverse prefixes than those that are allocated to ASes that do not transit traffic for others. As shown in Figure 1(c), roughly 97% of all prefixes announced by stub ASes (and more than 99% of all prefixes in the /24-/31 range announced by stub ASes) were announced from the same location.<sup>1</sup> The remain-

<sup>1</sup>Classifying an AS as a “stub” turns out to be difficult, as acquisitions, unorthodox transit relationships (*e.g.*, Harvard University appears as a transit for MIT in RouteViews), etc., preclude classifying the leaves of the RouteViews graph as stub ASes. Instead, we classify an AS as a stub



**Figure 1: Geographic diversity of IP addresses *within a single prefix*. Graphs show *complementary* CDFs for the *Clients* dataset; other datasets exhibit similar properties.**

ing prefixes announced by stub ASes, however, may contain locations that span large distances. For example, AS 6316 (StarNet) advertises a single /18 that contain hosts spanning over 2,000 miles in 9 locations. Another striking example is AS 4637 (Reach, an Asia-Pacific backbone “with direct connectivity to the US and Europe”), which advertises several /24 prefixes spanning over 10,000 miles (such as 202.84.142.0/24, which contains hosts in Perth, Australia and Dallas, Texas)!

About half of prefixes in the /8-/15 range contain IP addresses in multiple geographic locations, and about 97% of both prefixes longer than /24 and prefixes announced by stub ASes refer to IP addresses in only a single geographic location, which is expected. When stub ASes do advertise prefixes that contain hosts in different geographic locations, however, it is often the case that these hosts are not close together at all.

We hypothesized that, because large prefixes exhibit geographic diversity, large ASes might exhibit similar geographic diversity. That is, ASes with high degree (according to the RouteViews table) might announce prefixes from many diverse geographic locations. Interestingly, there are many small ASes that nevertheless announce geographically diverse prefixes as well: the correlation coefficient between AS degree and maximum distance between IP addresses contained within that AS is only 0.07, and many ASes with small degree commonly contain geographically diverse hosts. For example, AS 6509 (Canarie Inc., Canada), a relatively small organization with an out-degree of only 38 in the RouteViews table, announces a prefix 205.189.32.0/24 that spans locations that are 2,300 miles apart.

## 4.2 Discontiguous prefixes with single location

In this section, we analyze how frequently discontiguous prefixes (which cannot be aggregated) are announced by

if it has fewer than 5 downstream “customer” ASes per the classification algorithm from Gao [4].

Cause	Clients	Servers	Breadth
Fragmented Allocation	65.8	82.5	59.0
Load balance	1.5	1.9	3.9
Misclassification	4.5	4.8	13.8
Unknown	28.2	10.9	23.3

**Table 3: Analysis of the possible causes for the presence of discontiguous prefixes from the same geographic location within an AS.**

an AS from the same geographic location. We found that discontiguous prefixes formed between 70% and 74% of the total number of prefixes mapped in the three datasets. Discontiguous prefixes from the same geographic location and AS indicate that an IP prefix is too fine-grained.

Table 3 summarizes possible reasons for ASes announcing discontiguous prefixes from the same location, as well as their relative frequencies in our three datasets. Fragmented allocation is the single biggest reason for discontiguous prefixes being announced from the same AS and location: 65% of the discontiguous prefixes that appear in the routing table result from regional routing registries allocating discontiguous prefixes to ASes. We now analyze the causes for discontiguous prefixes in greater detail.

### 4.2.1 Fragmented allocation

IPv4 addresses are allocated by four Regional Internet Registries (RIRs): APNIC (Asia Pacific), ARIN (North America), LACNIC (South America and the Caribbean), and RIPE (Europe, Central Asia, and the Middle East).<sup>2</sup> The registries publish information on every block of IP space allocated by them. A typical allocation appears as:

```
arin|US|ipv4|19.0.0.0|16777216|19880615|assigned
```

This record specifies that a block of 16,777,216 contiguous addresses (*i.e.*, a /8) beginning from IP address 19.0.0.0, had been assigned to an organization on June 15th, 1988.

<sup>2</sup>In February 2005, a fifth RIR (AfrinIC) began full operation, covering registration for Africa. However, our datasets included the older registrations managed by ARIN and RIPE.

Registry	% fragment	% discontig	% all	% used
APNIC	25.11	31.90	30.97	81.07
ARIN	43.69	30.00	27.30	85.97
LACNIC	5.70	14.99	15.89	68.49
RIPENCC	25.50	23.11	25.85	86.38

**Table 4: Contribution of the various registries (*Breadth* dataset).**

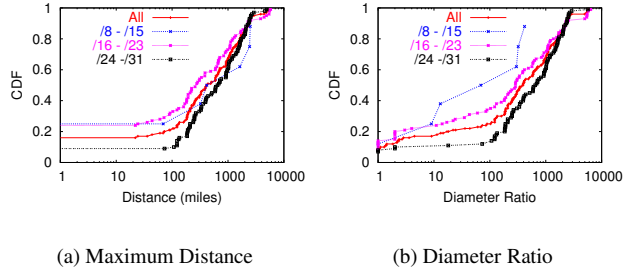
Using such allocation records, we investigated how often fragmented allocation was the cause for ASes announcing discontinuous prefixes. If a pair of discontinuous prefixes are from discontinuous allocations, then we conclude that an fragmented allocation has occurred.

Table 4 gives a registry-wise breakdown of the prefixes from fragmented allocations, discontinuous prefixes and the total number of prefixes observed. We have also tabulated the total fraction of the address space allocated at these registries. The table shows that LACNIC experiences less allocation pressure and similarly causes fewer fragmented allocations.

To further understand the reasons behind discontinuous allocations, we examined the allocation patterns of the 20  $\langle \text{AS}, \text{location} \rangle$  pairs in *Breadth* from which the largest number of discontinuous prefixes originated. We observed that 23% of the discontinuous allocations in these 20  $\langle \text{AS}, \text{location} \rangle$  pairs were made from discontinuous spaces *on the same day*, indicating that the registries were forced to make such assignments due to the paucity of IPv4 addresses. The remaining 77% of the allocations were made during different periods of time. Possible explanations for discontinuous address space allocations to an AS at different points of time are: (1) scarce IPv4 addresses are allocated conservatively to organizations, resulting in a fragmented set of addresses for each organization; and (2) two or more organizations with discontinuous addresses have one AS number due to a merger or acquisition.

#### 4.2.2 Load balance

An AS might announce a specific subnet of a bigger prefix in order to balance load over its two incoming links. For example, consider an AS with prefix  $p_i$  and two incoming links  $L_1$  and  $L_2$ , which desires that the traffic to a more specific (*i.e.*, “longer”) prefix  $p_j$  arrive through link  $L_1$  and the remaining traffic through link  $L_2$ . To achieve this goal, it announces the “longer” prefix  $p_j$  over link  $L_1$  and  $p_i$  over  $L_2$ . This practice is commonly referred to as “BGP hole punching”. Let  $D_{discontig}$  denote the set of all discontinuous prefixes in a dataset. To determine whether a pair of prefixes  $\{p_i, p_j\}$  appears in  $D_{discontig}$  due to hole punching, we check if their AS announces a supernet  $p_s$  that contains both  $p_i$  and  $p_j$  from the same location, thus producing a discontinuous pair of prefixes. We can observe from Table 3 that the number of discontinuous prefixes that appear due to load balancing is negligible—



**Figure 2: Geographic diversity of contiguous prefixes announced by the same AS. Graphs are for the *Breadth* dataset; other datasets show similar results.**

between 1.5% and 3.9% of the total number of discontinuous prefixes.

#### 4.2.3 Misclassification

As our location mapping data is incomplete, we could have misclassified a set of contiguous prefixes as discontinuous due to the absence of traceroutes to some prefixes. Consider a set of contiguous prefixes  $\{p_i, p_j, p_k\}$ . Assume that we have mapped  $p_i$  and  $p_k$  to a location  $L$ , but we do not have any location for prefix  $p_j$ . Then, by observing only prefixes  $p_i$  and  $p_k$ , we might mistakenly assume that the AS is announcing discontinuous prefixes from the same location. Hence, for every pair of discontinuous prefixes  $\{p_i, p_k\} \in D_{discontig}$ , we check if the “missing” intermediate prefixes are in fact announced by the AS in the RouteViews table. If so, we count this as an instance of misclassifying the pair  $\{p_i, p_k\}$  as discontinuous.

In Table 3, we observe that the *Breadth* dataset has more misclassifications than the other two. This result can be explained by the fact that, despite tracerouting to all advertised prefixes, we could not map all prefixes’ locations due to the limitations of `undns`. This limitation has a stronger influence on *Breadth* (which reached 161,974 prefixes) than on *Clients* (which reached 45,573).

#### 4.3 Contiguous prefixes with multiple locations

In this section, we study the extent to which ASes advertise contiguous IP prefixes that refer to networks in diverse geographic locations. We found 2,281 pairs of contiguous prefixes advertised by 384 different ASes. Of these pairs of prefixes, about one-fourth (607) of the pairs contained hosts in distinct geographic locations.<sup>3</sup> This finding suggests that the opportunities for aggregation may be less than that implied by the CIDR Report.

Figure 2(a) shows a CDF of the maximum distance spanned by hosts contained within a set of contiguous pre-

<sup>3</sup>Note that this measure is also a lower bound, as certain IP prefixes that we attributed to the same location might actually contain hosts in a different location that we did not probe.

fixes advertised by the same AS.<sup>4</sup> About 10% of all sets of contiguous prefixes were advertised from a single geographic location.

To better understand whether or not it makes sense to aggregate two contiguous prefixes, we defined a metric called the *diameter ratio* that highlights cases where a pair of contiguous prefixes represent two well-defined geographic clusters that are significantly far apart from each other. The *diameter ratio* is defined formally as follows:

$$\text{diameter ratio} = \frac{\text{maxdist}(L_1 \cup L_2)}{\min(\text{maxdist}(L_1), \text{maxdist}(L_2))}$$

where  $L_i$  is the set of locations contained in prefix  $p_i$  and *maxdist* is the maximum geographic distance between any pair of IP addresses in a set of IP addresses (*i.e.*, the “diameter” of the prefix). When either  $L_1$  or  $L_2$  contains only a single location, we set the denominator to 1. Intuitively, the diameter ratio is large when the locations within each of one or both of two prefixes are close together, but the aggregate set of locations are far apart from each other. A large diameter ratio may also reflect the case where the locations in one prefix are tightly clustered but the locations in the second are not. A large diameter ratio implies that aggregating the contiguous prefixes would remove the ability to express geographic routing policies.

Figure 2(b) shows the *diameter ratio* for each pair of contiguous prefixes in the routing table. We were surprised to see that smaller contiguous prefixes (*i.e.*, those in the /24-/31 range) spanned a greater geographic distance than larger contiguous prefixes (this phenomenon is shown in both Figure 2(a) and 2(b)). This geographic diversity is reflected along all three metrics (*i.e.*, number of distinct locations, maximum distance between IP addresses, and diameter ratio). Upon further examination, we found that this phenomenon can be explained by the fact that many ISPs based in the United States receive large prefix allocations and divide the allocation along /24 boundaries, advertising different /24s from different cities. On the other hand, we observe that ISPs in Europe and Asia typically advertise prefixes that correspond more closely with their actual allocations, which are usually considerably larger than /24. For example, in Europe, AS 5089 (NTL Group Limited, UK) advertises two separate contiguous /15s—80.2.0.0/15 and 80.4.0.0/15—for hosts in Cambridge and Luton, which are only about 75 miles apart.

To understand the extent to which the CIDR Report could be overestimating the opportunities for aggregation, we performed a CIDR Report style calculation on our dataset too. The CIDR Report computes the reduction in the number of contiguous prefixes when contiguous prefixes with same origin AS and AS path are aggregated. A similar calculation on our *Breadth* dataset showed that

the number of prefixes advertised can be reduced by 64% if we aggregate. However, aggregating geographically diverse prefixes could conflict with the traffic engineering goals of an AS. Hence, if we aggregate only the prefixes that in addition to having similar AS paths, are geographically “close” (we used diameter ratio  $\leq 500$  as a definition for “close”), then the number of announced prefixes could be reduced by only 20%. Thus, the CIDR Report could be overestimating the opportunities for aggregation by a factor of 3.

## 5. Conclusion

This paper studied the geographic properties of IP prefixes and their implications on Internet routing. Our findings have important implications not only for network applications that use IP prefixes to cluster end hosts, but also for Internet addressing. Advertising routes on a granularity that more closely reflects geographic locations (whether by renumbering, or by changing the addressing scheme entirely) could reduce routing table size by creating opportunities for aggregation.

## Acknowledgments

We thank Neil Spring for providing access to and support for *undns*. This research was conducted as part of the IRIS project (<http://project-iris.net/>), supported by the NSF under Cooperative Agreement No. ANI-0225660. Michael Freedman is supported by an NDSEG Fellowship. Nick Feamster is partially supported by an NSF Graduate Research Fellowship.

## REFERENCES

- [1] T. Bu, L. Gao, and D. Towsley. On characterizing BGP routing table growth. In *IEEE Global Internet Symposium*, Nov 2002.
- [2] CIDR report. <http://www.cidr-report.org/>, 2005.
- [3] M. J. Freedman, E. Freudenthal, and D. Mazières. Democratizing content publication with Coral. In *NSDI*, Mar 2004.
- [4] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. on Networking*, 9(6):733–745, Dec 2001.
- [5] G. Huston. Growth of the BGP table, 1994 to present. <http://bgp.potaroo.net/>, 2005.
- [6] B. Krishnamurthy and J. Wang. On Network-Aware Clustering of Web Clients. In *ACM SIGCOMM*, Aug 2000.
- [7] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang. IPv4 address allocation and the BGP routing table evolution. *ACM SIGCOMM CCR*, 35(1):71–80, 2005.
- [8] D. Meyer. University of Oregon RouteViews Project. <http://www.routeviews.org/>, 2005.
- [9] V. N. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for Internet hosts. In *ACM SIGCOMM*, Aug 2001.
- [10] PlanetLab. <http://www.planet-lab.org/>, 2005.
- [11] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *ACM SIGCOMM*, Aug 2003.

<sup>4</sup>When a set of contiguous prefixes had different mask lengths, we classified the prefixes according to the *minimum* mask length in the set.