

An Organization-Level View of the Internet and its Implications* (extended)

USC/ISI Technical Report ISI-TR-2012-679

June 2012

Xue Cai¹ John Heidemann¹ Balachander Krishnamurthy² Walter Willinger²

¹ USC/ISI, Marina del Rey, CA ² AT&T Labs Research, Florham Park, NJ
{xuecai,johnh}@isi.edu, {bala,walter}@research.att.com

ABSTRACT

We present a new clustering approach for mapping ASes to organizations, to develop an *organization*-level view of the Internet’s AS ecosystem. We demonstrate that the choice of clustering method and use of a new (though unconventional) data source in the form of company subsidiary information contained in the U.S. SEC Form 10-K filings are both essential to get accurate results. Evaluating our mapping and validating it against carefully chosen datasets shows few (less than 10%) false negatives for 90% of organizations and few false positives for 60% of our organizations. We apply our map to show the importance of an organization-level view of the Internet by contrasting it with the commonly-used view that considers only an organization’s “main” AS. We find that this main-AS view sometimes severely under-represents the influence of an organization in terms of announced addresses, geographic footprint, and peerings at Internet eXchange Points (IXPs). For example, for 20% of our organizations, the main-AS view detects only 10–60% of the cities covered by the corresponding organization-level view.

1. INTRODUCTION

Knowledge of and mapping the Internet’s connectivity structure is important to study network vulnerabilities to the various failure modes, be they technical [27, 28], political [20], or business-related [29, 38] in nature or the result of catastrophic events [9, 15] or intentional attacks [35]. However, by the Internet’s very design, there exist many shades of connectivity, from inherently physical links to types of virtual connections. To capture real-world aspects and be of practical relevance, abstractions of Internet topology must account for key features of this complex connectivity.

*Xue Cai and John Heidemann are partially supported by the US DHS, contract number D08PC75599, and by the NSF, grant number CNS-0626696. The conclusions of this work are those of the authors and do not necessarily reflect the views of DHS or NSF.

This work was begun when Xue Cai was an intern at AT&T Labs-Research.

A popular structure for studying the Internet has been the Internet’s AS graph [12, 14] where a node represents an *Autonomous System* (AS), commonly defined to be one or more networks in the Internet operated under a common routing policy [16], and links capture the exchange of reachability information among these ASes, each controlled by an *organization*. More than 30,000 ASes in today’s Internet appear in public BGP routing tables [26], and BGP-inferred AS-graphs have been extensively studied for more than a decade. One popular application has been the evaluation of Internet resilience, modeling threats to the network as removals of nodes or edges from the AS-graph [1, 10, 22, 23, 39], and measuring their impact in terms of network partitioning or increase in network diameter. However, modeling threats as graph operations can be problematic when the graphs fail to account for lower- or higher-level real-world details, such as geographically distributed routers and physical connections or organizational alliances.

Moving beyond the Internet’s physical infrastructure, many issues concerning today’s network occur at the *organization* level, above the AS-level, and arise from business or political disputes. Examining these issues when organizations operate multiple ASes requires understanding the *AS ecosystem* [6] that aims to reflect organization-level constraints imposed on the traditional AS-graph.

The main contribution of this paper is to evaluate the effect of this organization-level structure on important aspects related to Internet topology. To this end, we show that multi-AS organizations matter in today’s Internet: some 36% of assigned AS numbers and 29% of actively routed ASes belong to multi-AS organizations. More importantly, organizations using multiple ASes in routing are particularly prominent, announcing nearly two-thirds of all routed addresses. Using historic routing table snapshots, we analyze underlying causes of multi-AS usage and illustrate that multi-AS organizations are not a transient phenomenon.

We evaluate some effects of this organization-level

structure on the Internet topology. In particular, we show that the traditional view that typically considers an organization’s “main” or “best-known” AS greatly underestimates the geographic footprint and IP address coverage when compared to an organization-wide view that encompasses all of an organization’s routed ASes. We also demonstrate that understanding the public peering of companies is strengthened by an organization-level view.

These examples illustrate how our work provides a basis for gaining a deeper understanding of the economic relationships in today’s Internet inside an organization and between organizations. This need is growing: We expect these relationships to continue to evolve as the Internet changes, particularly as business models become more heterogeneous (as ISPs distribute content, and content providers deploy networks).

Consideration of organization effects is only possible with adequate organization-level maps of the Internet. This paper develops a new clustering algorithm that add new data source for mapping ASes to companies, company subsidiary information contained in the U.S. SEC Form 10-K filings. With a new approach to hierarchical clustering with weighted attributions, this approach provides much greater accuracy than our prior work’s much simpler clustering without 10-K information [6]. Given the legal requirements that form the basis of this newly used information source, its accuracy is superior to the largely volunteer-based efforts such a Packet Clearing House [24] or the Regional Internet Registries (RIRs) [25]. We show that, especially for U.S.-based companies, the value of the 10-K data can greatly improve organization completeness. Although the use of 10-K data is not without problems, we circumvent inconsistency in the names of companies and its subsidiaries through simple heuristics.

We evaluate our new approach to mapping ASes to organization as thoroughly as possible, using four datasets composed of more than 100 organizations and 4,000 ASes. These datasets were chosen to balance quality, unbiasedness, and size. We show that our results are accurate, with 90% of the organizations showing false-positive rates of less than 10%; they are also quite complete with more than 60% of the organizations showing 10% or fewer false-negatives. Our evaluation shows how data quality, data availability, and the type of clustering algorithms affect our results. We find that company subsidiary information from U.S. SEC Form 10-K filings significantly reduce false negatives.

To improve research using Internet topologies and allow others to build on our work, our current AS-to-organization maps will be released in June 2012, along with our public test data and source code.

Beyond the basic paper, this technical report adds several appendices with additional supporting informa-

tion. Appendix B provides more information on how we select and normalize the AS attributes used for clustering. Appendix C describes in more detail how we automatically train weights for the different attribute types. Appendix D discusses specifics on how we manually verify and use 10-K links. Appendix E evaluates a volunteer-based AS-to-Organization mapping provided by Packet Clearing House (PCH) and compares it to our mapping effort. Lastly, Appendix F describes an analysis of the persistence of multi-AS usage.

2. RELATED WORK

While there exists an extensive and still growing literature on various aspects of measuring, modeling, generating, and analyzing Internet connectivity structures at different levels (e.g., router-level, AS-level, overlays like the Web, P2P, or online social networks), little attention has been paid to date on organization-level effects that capture the constraints imposed on those ASes that operate as part of multi-AS organizations or companies. Our work seeks to broaden the existing AS perspective by exploring these organization-level effects.

To the best of our knowledge, only two prior groups have considered how ASes relate to organizations. First, to study observed incongruities between AS paths derived from traceroute probes and BGP routing data, Hyun et al. [17] inferred AS ownership based on the ID and name of registered owners (organizations) in a subset of the American Registry for Internet Numbers (ARIN) WHOIS database relying on a largely manual process. In this paper, we not only demonstrate that a more complete and diverse input dataset is absolutely critical for successfully mapping ASes to their organizations, but also present a fully automated process for performing this AS-to-organization mapping. Second, Packet Clearing House (PCH) [24] maintains a manually generated AS/organization directory that relies on voluntary efforts and is hosted by PCH for the main purpose of facilitating contacts between different network operators. We evaluate the information contained in this directory in our technical report [7] where we comment on the shortcomings inherent in data resulting from voluntary efforts as compared to information that is provided, say, for legal reasons.

3. METHODOLOGY

We map ASes to organizations through a combination of two methods: automatic clustering done on a structured data source (Section 3.1) and a semi-automatic method on a less structured data source (Section 3.2).

3.1 Automated Clustering with WHOIS Data

Our automatic method relies on publicly available information from AS registration data in WHOIS (Section 3.1.1), and consists of the following four steps that

	All	OrgID	Contact	Phone	Email
ARIN	22k (100%)	21k (95%)	20k (91%)	19k (86%)	19k (86%)
RIPE	20k (100%)	13k (65%)	19k (95%)	<i>unavail.</i>	14k (70%)
APNIC	6k (100%)	<i>unavail.</i>	6k (95%)	5k (83%)	5k (83%)
LACNIC	1.5k (100%)	1.5k (100%)	<i>unavail.</i>	<i>unavail.</i>	<i>unavail.</i>
AfriNIC	0.6k (100%)	0.6k (87%)	0.6k (100%)	0.6k (98%)	<i>unavail.</i>
<i>All</i>	<i>50k (100%)</i>	<i>36k (72%)</i>	<i>46k (92%)</i>	<i>25k (50%)</i>	<i>38k (76%)</i>

Table 1: Data availability (AS count) for four attribute types across the 5 RIRs.

are discussed in more detail in the rest of this section:

1. Extract, standardize and link attributes to each AS (Section 3.1.2).
2. Train weights of attribute types to reflect different degrees of importance (Section 3.1.3).
3. Compute similarity score between each AS pairs based on their weighted attributes (Section 3.1.4).
4. Cluster ASes based on the similarity scores computed in step 3 and label each cluster (Section 3.1.5).

3.1.1 WHOIS Data

The *WHOIS* database stores AS-specific registration information that is provided by each AS on a voluntary basis. *WHOIS* was originated to assist network operators in contacting each other, but there are no forcing mechanisms in place to ensure that the information provided by each AS is complete or accurate. Each of the five Regional Internet Registries [25] (RIRs) provides its portion of *WHOIS* data, with ARIN using its own format [2] and the other RIRs relying on the Routing Policy Specification Language (RPSL) [30]. To make full use of this *WHOIS* data, we first merge these different formats into a common format.

WHOIS data is composed of various types of records, and each record is associated with multiple attributes. We are interested in three general types of records in *WHOIS*: *Autonomous Systems* (ASes), *organizations* (or orgs), and *points-of-contact* (contacts). From specializations of these types (for example, administrative or technical contacts), we identify a total of 66 different types of AS attributes that provide potentially useful information for identifying ASes with an organization. In particular, there are some 50k ASes in all five RIRs (see Table 1), identified by *ASHandle* records in ARIN and *aut-num* records in other RIRs. Some AS’s org records are linked by *OrgID* or *org* attributes in their records. Org records are often used in *WHOIS* for common management of multiple resources and are potentially useful for identifying an AS’s organization. However, RIR policies do not require a one-to-one mapping between *WHOIS*-derived and real-world organizations, making simple clustering on common organization records ineffective. In terms of an AS’s contact records, they typically identify individuals in charge of administrative, technical, abuse or operations aspects of the AS.

For example, while a multi-AS organization may use the same contact information for all its ASes, even if it uses different contacts, they can often be linked based on common telephone numbers and e-mail addresses.

We pursue two strategies with respect to the above-mentioned attributes. The first strategy is to aggregate them into four attribute types, namely *OrgID*, contact ID, phone, and email. The second strategy is to keep the 66 identified attributes separate so as to allow for a differentiation between their subtle semantic meanings. For example, by treating administrative, technical, abuse and network operation center (NOC) phone and e-mail contact information as separate attributes, we account for the possibility that different types of contacts have different importance and thus should be treated differently. To illustrate, an administrative contact is more likely to be a parent organization employee than an employee of a third-party outsourcing company; however, just the opposite may be the case with respect to a technical contact. Throughout this paper, we refer to these two sets of attributes as *4attr* and *66attr*, respectively.

Challenges: There are several challenges in using *WHOIS* data to map ASes to organizations. A major problem arises from the fact that the *WHOIS* database is incomplete, and the records can be stale and incorrect. Table 1 shows the data availability of four different attribute types. Additionally, we cannot associate any single type of attribute (*OrgID*, contact ID, phone and email) for all ASes, and some RIRs filter all information due to privacy concerns (for example, RIPE filters all phone numbers in bulk data). These challenges require our use of clustering across multiple attributes.

Second, the use of outsourcing companies complicates use of *WHOIS* data, resulting in incorrectly clustering unrelated organizations as one. When an organization has a third-party handle network operations, the contact information present in an AS record may not link to an employee of the AS’s parent organization, but instead identify an employee of the third-party outsourcing company. We discuss these cases and our solutions in detail in Section 4.4.

Finally, mergers and acquisitions are primary reasons for mismatches between real-world organizations and *WHOIS*-based organizational identities. If *WHOIS* records for ASes of an acquisition are not updated to refer to the new parent company, or if the acquisition maintains distinct AS-level administration (for example, as Youtube and Google), then it may be impossible to infer the correct organizational identity of that AS from *WHOIS* information alone. In Section 3.2 we turn to additional data to explicitly identify acquisitions.

3.1.2 Attribute Extraction and Standardization

The design of any clustering algorithm must select

and normalize the features used for clustering - in our case, ASes and their attributes. Ideal attributes will link all ASes of one organization, and will not link them with other organizations. However, since no single attribute in WHOIS meets these ideals, we use a combination of attributes to maximize accuracy. To this end, we first extract raw attributes from *WHOIS*, clean them by canonicalizing them to simple attributes, and finally discard generic attributes (e.g., Gmail and Hotmail for e-mail addresses) to produce *4attr* and *66attr sets*. (See Appendix B for details.)

3.1.3 Training Attribute Weights

Attributes have different degrees of importance based on their types. To account for this fact, we assign to the different attributes weights that have been tuned based on training data. Below, we briefly summarize our approach. More details can be found in [7].

To train the attribute weights, we use parallel hill-climbing [32] over a training set of about 10,000 ASes. This training set consists of some 715 ASes for which we have reliable organizational identity information and 9k additional ASes (“noise”). The algorithm aims to minimize the likelihood of mis-clustering (i.e., assigning an AS to the wrong cluster, a false-positive) or missing to assign an AS to the correct cluster (false-negative) those ASes whose organizational identity we know. Specifically, to avoid having outliers unduly influence training, we use as our objective function the sum of the quartiles of the false-positive and false-negative rates over those ASes.

The size of training set (10,000 ASes) is carefully selected based on time and hardware constraints. Clustering with 50k ASes (whole population) takes about 3 days and 24GB memory. Because we have limited access to such large-memory hardware, we use this training set to study parameter effects, so that each round requires only 20 minutes and 1GB of memory. In all we examined about 15k weight vectors, using about 200 days of compute time, done in parallel over about a week.

Clustering shows a best weighting scheme with $\hat{w}=\{0.75, 0.1, 0.1, 0.05\}$ for the four attribute types (OrgID, contact ID, phone and e-mail) in the *4attr* set. It emphasizes the correctness of OrgID (0.75), downplays contact ID (0.1), phone (0.1) and email (0.05). OrgIDs are intended for common administrative management thus are unlikely to cause false positives. Contact IDs, phones and emails, though, can be registered by outsourcing third parties thus have the potential to introduce false positives.

Weighting schemes for the *66attr* set are generally worse (1.5 to 2 times larger objective function), so we rule out using 66 attribute types in practice. We examine causes for this worse performance, and find more

attributes dividing a general attribute type into many specific types, often breaking clustering links and thus leading to higher false-negative rates. Consider, for example, the e-mail attribute type (similar arguments apply to phone and contact ID). With *4attr* input, administrative and technical contacts for two ASes `admin@as1.example.com` and `tech@as2.example.com` are part of the same attribute and so link the ASes. However, with *66attr* input, these attributes are considered independent and so will fail to link the ASes.

3.1.4 Similarity Matrix

Next we use the weights assigned to the different attribute types to link ASes based on a *similarity score*. We compute similarity scores for all AS pairs and store them in a *similarity matrix*. This matrix is used as input in Section 3.1.5 to cluster ASes.

We use *weights* and *Jaccard index* to compute the similarity score between two ASes. Let $s_{x,y}$ denotes the similarity score between AS x and y , \hat{w} the weight vector of which w_i , $i \in [1..M]$, $\sum_{i=1}^M w_i = 1$, is the weight of the i^{th} attribute type, and \hat{X} denote the attribute vector of AS x of which X_i , $i \in [1..M]$, is the attribute set of the i^{th} attribute type. Then we have

$$s_{x,y} = \sum_{i=1}^M w_i \cdot J(X_i, Y_i) \quad (1)$$

where the Jaccard index $J(X_i, Y_i)$ is the similarity score between two specific sets of attributes of the same type (the i^{th} attribute type); e.g., between email attribute set `{comcast.com, comcast.net}` and email attribute set `{comcast.com}`, and defined as

$$J(X_i, Y_i) = \frac{X_i \cap Y_i}{X_i \cup Y_i}$$

Note the way we calculate similarity score implicitly suggests that different types of attributes are *orthogonal*. Only attributes of the same type are compared. With *4attr* input ($M = 4$), all e-mails are compared with each other; similar for OrgIDs, contact IDs, or phone numbers. However, with *66attr* input ($M = 66$), since administrative contact and technical contact are treated differently, their IDs, e-mails and phone numbers are of different types, thus an administrative e-mail will never be compared with a technical e-mail.

3.1.5 Clustering Algorithm

The final stage consists of using the similarity matrix to cluster ASes so that ideally, each generated cluster corresponds to a real-world organization and can be labeled appropriately based on clues such as domain names or keywords.

To cluster ASes, we rely on a *hierarchical clustering* algorithm. It starts with a set of individual ASes with their pairwise similarity matrix. In the beginning, each

individual denotes a cluster. During each round, two of the *closest* clusters get merged together until there is only one cluster remaining. The algorithm produces a hierarchy (or binary tree) of clusters and it is the user’s choice to decide at which level to cut the tree. The higher the level, the smaller the similarity between the clusters at that level, and the number of clusters is lower as well. We decide the similarity threshold that cuts clusters by the following automatic training method. Because similarity score is proportional to the sum of weights defined in Equation (1), during weight training, we set similarity threshold to be a fixed value first, then let the training algorithm walk through different weight vectors. After the best weight vector is selected, we normalize the scheme and adjust the similarity threshold based on this normalization.

There are several ways to define the distance between two clusters during clustering, and the choice of distance can significantly affect the clustering results. The four definitions are *maximum-linkage* (the maximum distance among all pair-wise elements), *single-linkage* (the minimum distance), *average-linkage* (the average distance), and *centroid-linkage* (the distance between the centroids of two clusters). We reject maximum-linkage as too strict, creating false negatives, and single-linkage is too aggressive, causing too many false positives, and we avoid centroid-linkage as too computation intensive. We choose average-linkage because it can best help us explore the underlying semantics of the *WHOIS* dataset (see Section 4.4) and deal with some of the challenges posed by it (e.g., tech-outsourcing issues).

3.1.6 Cluster Labeling and Selection

Hierarchical average-linkage clustering produces a set of AS clusters which, in turn, need to be labeled with information that identify the corresponding organizations. One promising information source for labeling is the email domain used for clustering. It is usually human-readable and can provide accurate reference to the organizations’ websites. Other attribute types can be obscure (e.g., OrgIDs with unusual abbreviations), or may leak private information (e.g., contact IDs and phones). We also extract text names in AS and OrgID records, including *ASName* or *as-name*, *OrgName*, *descr* (description), and *owner* fields. To improve search speed and label quality, we break these names into keywords and rank them based on their frequency and uniqueness. This process tends to highlight the keywords that help in identifying an organization’s identity.

To manually identify an organization’s AS cluster, one can search for related keywords or domains, and manually decide which AS cluster is the closest. In contrast, to automatically identify all organizations’ AS

clusters, one needs a list of keywords or domains and a judge function to pick one or multiple matched AS clusters. However, we do not have such list of keywords or domains for every Internet-related organization. Thus, when comparing the accuracy of our results with the ground truth, for the sake of simplicity, we always pick the *biggest* cluster in our results to compare with the ground truth cluster. We caution that this decision favors both lower false-negative rate and higher false-positive rate.

3.2 Semi-automatic Clustering with 10-K Data

A particular problem we encounter using *WHOIS* data to mapping ASes to their organizations is mergers and acquisitions. Such changes often result in stale information, reducing mapping accuracy. To address this issue, we advocate in this paper the use of a new and previously untapped source of information for AS-to-organization mapping - company subsidiary data contained in the U.S. SEC Form 10-K filings.

3.2.1 Form 10-K Data

Unlike the voluntary nature of *WHOIS*, 10-K forms are mandated by law and are required to be filed annually by all publicly-traded U.S.-based companies. Because of this legal requirement, the completeness and accuracy of the data is superior to the voluntary *WHOIS* database. Form 10-K therefore represents the ground truth for subsidiary relationships among publicly-traded U.S.-based companies for the year prior to the filings. It’s disadvantages are that it does not apply outside the U.S., and 10-K names can be imprecise as we describe below.

Form 10-K data is freely available from the U.S. Securities and Exchange Commission’s (SEC) EDGAR database [34]. EDGAR covers each of the thousands of publicly-traded, U.S.-based companies. Each form is identified by a unique company identifier (we call the *10-K ID*), the year the form is filed, and a list of all of a company’s current subsidiaries. After extracting the subsidiary names from all these lists and normalizing them to lower case, we associate the company name present in the beginning of each 10-K form and all subsidiaries with the 10-K ID. In total, we extract 8,706 companies and 156,936 subsidiaries from the database for fiscal year 2010.

The main weakness of 10-K data is that entries are imprecise because company names are not standardized across 10-K and *WHOIS*. Fully normalizing names is a difficult problem in natural language understanding, mainly because many variations are context dependent and some names provide very little context. For example, *Network* and *Inc.* are two words that convey in general little information about organization identity, but they cause noise and cannot be dropped for

matching. Similarly, variations in spacing, abbreviations, and level of detail can all cause errors in name matching, and some manual data cleaning is necessary to determine that, for example, “Apple”, “Apple Computer”, and “Apple, Inc.” all refer to one and the same company, while “Apple Records” identifies a different company.

3.2.2 Automatic Name Linkage

Given that the Form 10-K data is an accurate source of company subsidiary information, in theory we can use this data to cluster ASes that belong to different subsidiaries of one and the same organization. In particular, since Form 10-K data provides the names of a company’s subsidiaries and WHOIS data includes the name of each AS, by simply matching these names, we can cluster apparently different subsidiaries. Among the challenges we face in practice is that mergers or acquisitions often result in subsidiaries that retain a distinct identity as indicated by their names or WHOIS entries. For example, Google maintains a public face for Youtube, but not for Postini, a later acquisition. Moreover, the names are in general not normalized, and may thus provide little context, resulting in name matching that is error-prone.

To overcome these difficulties, we first describe an automated procedure for linking AS names with names of subsidiaries in this section, and then detail a manual process for verifying and pruning links for some 50 purposefully chosen organizations in Section 3.2.3.

Linking the same entity from different data sources, usually with different names, is called *record linkage*, and is a well studied problem in data mining area. A number of existing algorithms [8,31,33] are readily available, and for the problem at hand, we choose the TF-IDF method [33], mainly because it pays special attention to infrequent keywords in names and is fairly straightforward to implement without complicated parameter configuration. TF-IDF is commonly used to query natural text documents in corpus by certain keywords. While we believe it is a reasonable choice for our application, it is not optimized in any way for the problem at hand; that is, matching company names.

3.2.3 Manual Verification and Pruning

An error in name matching can result in falsely clustering hundreds of ASes to a single organization. Since these errors stem from automatic clustering based on matches of imprecise names, rather than using all automatic name matches, we manually verify and use links for some 50 purposefully selected organizations.

We select 50 organizations (about 0.6% of the 8,706 organizations) intentionally to favor those most important in the real world and to the Internet’s ecosystem. We select 38 large, computer-related organizations from

the Fortune 500, and add 12 additional large ISPs. Of the 1817 links that the automated clustering produced for these 50 organizations, we verify and keep 1226 links, dropping 591. Details about these organizations and pruning are in Appendix D

3.2.4 Enhanced AS Clustering

To make good use of the *Form 10-K* data, we modify the clustering algorithm described in Section 3.1.5 by incorporating the information contained in the above-identified links between ASes and 10-K organizations.

To this end, we augment the $4attr$ set with new attributes representing the 10-K organization an AS is linked with (as determined in Section 3.2.3 above), and call these new attributes *10-K ID* attributes. Together with previous WHOIS attributes in four types, this augmented set now comprises attributes in five types (OrgID, contactID, phone, e-mail and 10-K ID), and is referred as $4attr+10K$. We assign 10-K ID attributes the same weight as OrgID attributes; we use a strong weight because we have manually verified the accuracy of these new attributes. The new weight vector $\hat{w}' = \{0.75, 0.1, 0.1, 0.05, 0.75\}$ is not normalized and we keep the same similarity threshold when cutting the clustering tree (see Section 3.1.5). As a result, on one hand, if two ASes are linked with the same 10-K organization, their similarity score will be higher than before, and thus more likely to be clustered together later; on the other hand, if two ASes are linked only by WHOIS attributes, their similarity score can stay the same and whether they end up in the same cluster or not will be the same as that before using 10-K information.

We re-computed the similarity matrix with the $4attr+10K$ set as input and fed the result to the clustering algorithm described in Section 3.1.5. We compare the clustering results obtained by this $4attr+10K$ -based method with those produced by the $4attr$ -based approach in the subsequent section.

4. VALIDATION

To validate the accuracy of our clustering methods, we use four datasets (Table 2 and Section 4.1), purposefully chosen to trade-off confidence among three criteria: quality, unbiasedness, and size of ground truth. We describe the validation method in Section 4.2 and present our results in Section 4.3, focusing in particular on how specific aspects of our methodology and properties of our datasets impact accuracy (Section 4.4) and lead to significant improvements in accuracy.

4.1 Validation Datasets

In order to gain a comprehensive view of the quality of our mapping results, we use four different datasets with varying degrees of quality, unbiasedness and size for evaluation. Table 2 lists them and ranks them by

Datasets	Quality	Unbiasedness	Size		
			Orgs.	ASes	
T_{tier1}	operator-provided (nearly definitive)	intentionally selected (potentially biased)	1	many	small
T_{9org}	records-inferred, abundant records (very good)	intentionally selected (potentially biased)	9	502	medium
$T_{randtop}$	records-inferred, mostly abundant (good)	randomly selected from top (mostly unbiased)	50	2516	large
$T_{randall}$	records-inferred, few records (potentially incomplete)	randomly selected from all (completely unbiased)	50	1001	large

Table 2: Validation datasets ranked by quality, unbiasedness and size.

these criteria.

The first dataset, T_{tier1} , is provided by a Tier-1 operator, and thus represents the highest quality in terms of completeness and accuracy. However, containing only one organization, it is clearly biased towards that organization.

Obtaining operator-provided ground truth is difficult, so we next infer three datasets from public records. From public online documents, routing data and *WHOIS* information, we believe we find most ASes of a given organization for our targets (all public companies).

These three datasets consist of different samples of organizations, for different purposes. T_{9org} contains nine big U.S.-based public companies with plenty of information online, thus it is of fairly good quality. Although hand-picked, it sheds light over key players in today’s Internet, that is, four large telecommunications companies, four content providers, and a root-DNS provider.

In contrast, $T_{randtop}$ and $T_{randall}$ are randomly chosen, and each consists of 50 organizations. We first consider all clusters that were produced by our clustering method that uses *4attr+10K* as input, take a random sample, and finally identify the organization identity of the sample from their AS *WHOIS* records. More precisely, $T_{randtop}$ is a random sample of size 50 from the 100 largest organizations we find, where the size of an organization is given by the number its ASes. From manual inspection, this dataset contains large ISPs, big research networks, media conglomerates and multi-national financial companies. By comparison, $T_{randall}$ is a randomly selected set of 50 organizations from all 36,463 clusters our method produces. Most of $T_{randall}$ are small, private organizations, often without even a website. To summarize, although slightly less complete than T_{9org} (because of less public information), $T_{randtop}$ sheds light over a broader and less biased range of key Internet players. Although mostly comprising small and less interesting organizations, $T_{randall}$ represents a completely unbiased sample to evaluate our algorithm’s accuracy.

There are six organizations in common between $T_{randtop}$ and T_{9org} , while $T_{randall}$ and $T_{randtop}$ are disjoint. Also, while the median organization size in $T_{randall}$ is 1, overall it includes 1001 ASes because one organization (a network information center) has 944 ASes. When ignoring those ASes, we end up with 49 organizations

with a total of 57 ASes.

4.2 Validation Method

To validate our results, we consider the clusters produced by our clustering method that takes *4attr+10K* as input. For each organization in our validation sets, we first select the biggest cluster (see Section 3.1.5) that overlaps with the ground truth and compare them. We then check how many ASes are wrongly assigned to the cluster (false negatives) and how many ASes are missing from the cluster (false positives). We define false positives (*fp*) and false negatives (*fn*) as follows. Let M_i be the i^{th} cluster in the ground truth (e.g., T_{9org}) and C the biggest cluster in our results overlapping with the ground truth. Then we have

$$fn = 1 - \frac{|M_i \cap C|}{|M_i|}, fp = \frac{|C|}{|M_i|} - \frac{|M_i \cap C|}{|M_i|} \quad (2)$$

where $C \in R_{ours}$ is a cluster produced by our method and M_i is the cluster in the ground truth overlapping with C (e.g., $M_i \in T_{9org}$).

For simplicity, we also classify our validation results for each organization into *good* or *bad* with the help of the false-positive and false-negative rates. In particular, we call the results for an organization *good* if both types of mistakes are below 10% and *bad* otherwise.

4.3 Validation Results

In this section, we first present the overall findings for all four validation datasets (Section 4.3.1). We then look into the underlying causes of mistakes and the obstacles we face when mapping ASes into organizations. (Section 4.3.2).

4.3.1 Overall statistics

Table 3 summarizes the results for T_{tier1} and T_{9org} , showing overall very low false-positive rates and moderately low false-negative rates (see Section 4.3.2 for causes). As we can see from the *false-positive analysis* shown in the top portion of the table, for nine out of ten organizations (90%), we wrongly cluster less than 10% of ASes. In terms of the *false-negative analysis* given in the bottom portion of the table, we find all ASes for six organizations (60%), more than 80% for two other organizations, and perform poorly for the two remaining ones. The false-positive and false-negative rates for T_{tier1} alone are 7% and 27%, respectively.

category	orgs	percentage
<i>false positive</i>		
good	9	90%
perfect (=0%)	6	60%
0%-10%	3	30%
bad	1	10%
10%-20%	1	10%
<i>false negative</i>		
good	6	60%
perfect (=0%)	6	6%
bad	4	40%
10%-20%	2	20%
20%-40%	2	20%

Table 3: Validation of 10 intentionally selected organizations including a Tier-1 ISP.

category	orgs	percentage
<i>false positive</i>		
good	47	94%
perfect (=0%)	31	62%
0%-10%	16	32%
bad	3	6%
10%-20%	2	4%
20%-40%	1	2%
<i>false negative</i>		
good	34	68%
perfect (=0%)	23	46%
0%-10%	11	22%
bad	16	32%
10%-20%	6	12%
20%-40%	8	16%
>40%	2	4%

Table 4: Validation of randomly selected organizations from top 100 clusters.

Similar results hold for $T_{randtop}$. Table 4 shows that 47 organizations (94%) have fewer than 10% false positives and that we found more than 90% ASes for 34 organizations (68%). Not only do these numbers show that our results generalize to large organizations, but they also confirm that our weights are not overfitted as a result of using the ten organizations in T_{tier1} and T_{9org} for training.

Finally, validation based on the truly unbiased dataset $T_{randall}$ shows that our method performs even better for the majority of Internet-related organizations. In fact, as shown in Table 5, for almost all organizations (48, or 96%), the results are good with respect to false positives, and for almost as many (47, or 94%) are good for false negatives. The high accuracy with respect to false-negatives follows from the fact that the majority of Internet-related organizations are simple and small, which makes finding all their ASes easy; indeed, 43 out of the 50 organizations in $T_{randall}$ have only one AS.

4.3.2 Understanding Sources of False-Positives and False-Negatives

While the overall accuracy of our results is quite good, our approach does poorly in some cases. We next ex-

category	orgs	percentage
<i>false positive</i>		
good	48	96%
perfect (=0%)	48	96%
bad	2	4%
>40%	2	4%
<i>false negative</i>		
good	47	94%
perfect (=0%)	47	94%
bad	3	6%
>40%	3	6%

Table 5: Validation of randomly selected organizations from all clusters.

amine these cases to understand the limitations of our approach and suggest possible future improvements.

False Positives: The main cause for false positives is the lack of a clear boundary between organizations. A typical real-world scenario involves ISPs and IT consulting companies that often provide technical support, including the management of AS records, for their customers. Thus, they share the same contact information with their customers which, in turn, becomes a common reason for false-positives in our results. This scenario applies to the three organizations with bad false-positive rates in Table 4, two ISPs and one tech-outsourcing company, also to the two organizations in Table 5 with more than 40% false-positives.

False Negatives: The main cause for false-negatives is missing or inaccurate subsidiaries. Although we use company subsidiary information via the $4attr+10K$ input set to our clustering algorithm, we encounter numerous cases where different subsidiaries maintain their distinct identities in *WHOIS* (such as the two organizations with bad false-negative rates in Table 4, Nippon Telegraph and Telephone (NTT) and Deutsche Telekom, which are both foreign and large). *WHOIS* records that are out-of-date and do not reflect the correct subsidiary names make it difficult to produce accurate clusters without external knowledge and result in false-negatives in our approach. These difficulties make false-negative rates bad for the two organizations in Table 4. They are Nippon Telegraph and Telephone (NTT) and Deutsche Telekom, which are both complex and non-U.S. companies.

4.4 Factors that Improve Accuracy

Figure 1 shows the significant improvements that our new method (*new*) is able to achieve when compared to our earlier approach (*old* [6]). In terms of false-negatives (the green arrows annotated by “!”), the results for all ten organizations improved, from 2% to 36%. In addition, four organizations (i.e., the Tier-1 ISP, Verizon, Limelight and ISC) show improvements with respect to their false-positive rates, from 4% to

129%. Only Akamai shows a slightly worse false-positive rate, at 3%. Manual inspection shows that this change is due to a new outsourcing arrangement.

We next evaluate what specific aspects of our new algorithm helped improve the accuracy of our results and highlight the impact that our new data source (i.e., *Form 10-K* data) has on AS clustering.

4.4.1 Avoiding incorrect assertions (false-positives)

A good clustering method should be able to clearly identify organizational boundaries and put each AS into its own and relevant cluster. However, organization boundaries become usually blurred by a number of real-world organizational relationships, including tech-outsourcing, technical support, and joint ventures. These blurred organizational boundaries result in false-positives when we try to associate ASes with their organization.

To sharpen our discovered organizational boundaries and thus reduce false-positives, we exploit hierarchical average-linkage clustering (see Section 3). Unlike clustering where any linkage joins two clusters, average-linkage examines all links between any two ASes in two clusters to judge the strength of the relationship. This weighting results in tenuous links between otherwise well connected clusters getting severed, preventing organizational boundaries from blurring together.

Among the organizations whose false-positive rates improved significantly, ISC and Verizon stand out. ISC’s false-positive rate is drastically reduced from 129% to 0% and Verizon’s is reduced from 28% to 3%. We confirm that 70 of the 71 false-positives for ISC were caused by a single linkage that existed because of tech-outsourcing. For Verizon, 59 of the original 66 false-positives were caused by several single linkages as a result of a combination of customer and tech-outsourcing relationships. In both cases, the use of average-linkage results in less aggressive clustering by ignoring these weak relationships.

4.4.2 Factors affecting completeness (false-negatives)

Besides clearly identifying organizational boundaries, a good clustering method should also make use of all critical information that relates an AS to its organization. For example, contact information typically relates an AS to its direct operator. However, if this direct operator is a subsidiary of a big company, relating it to its parent company may require additional information. Unavailable critical information can result in missing ASes, which in turn makes our results incomplete.

To reduce false-negatives and mitigate the problems caused by subsidiaries, our new algorithm relies on a combination of more complete and up-to-date *WHOIS* data and a novel information source in the form of *Form 10-K* data. In particular, we note that two organizations (i.e., Akamai (2 ASes) and CN Mobile (1 AS))

category	orgs	percentage
company subsidiary information used	18	100%
big improvement (>20%)	4	22%
medium improvement (10-20%)	2	11%
small improvement ($\leq 10\%$)	8	44%
no improvement (=0%)	4	22%

Table 6: Improvement on false-negative rate when company subsidiary information is used.

benefited from using more attributes than in [6]. An up-to-date *WHOIS* database is also very important as can be seen in the cases of Yahoo (4 ASes), Google (3 ASes), ISC (2 ASes) and Limelight (1 AS). We next examine in more detail the effect that the use of *Form 10-K* data had on our results.

4.4.3 Does company subsidiary information help?

We next explore the role of 10-K data in improving accuracy. To examine its benefits, we consider *Form 10-K* data for 50 U.S.-based public companies and compare two sets of clustering results. One set was obtained by running our new clustering algorithm with *4attr+10K* as input (i.e., using company subsidiary information), the other running the same algorithm with *4attr* as input. We evaluate both sets of results using all four validation datasets and check for improvements concerning the false-negatives.

Of the 50 companies considered, 18 are in our ground truth datasets. As shown in Table 6, 14 out of 18 organizations improve. The improvements are significant for four organization (i.e., Limelight, Oracle, IBM and HP) because this new source of information captures most of their subsidiaries. It also helps four organizations (i.e., Cogent, VeriSign, Yahoo, Comcast) where we capture all of their ASes, obtaining a 0% false-negative rate. In short, although our approach can miss ASes for organizations that tend to have a complex structure or long history of mergers and acquisitions, company subsidiary information in the form of publicly available *Form 10-K* data is able to alleviate this problem, resulting in often significant improved AS clustering.

5. PROPERTIES OF THE ORGANIZATION-LEVEL INTERNET

In this section we study our AS-to-organization mapping results, particularly for organizations that use multiple ASes. We first examine the *prevalence* and *influence* of multi-AS usage in Section 5.1. We then investigate why organizations use multiple ASes in Section 5.2.

5.1 Relevance of multi-AS organizations

To quantify the relevance of multi-AS organizations, we analyze our AS-to-Org mapping results and show in Table 7 that 49,262 ASes map into 36,463 organizations in total (top row), of which 13% (4,856) are *multi-*

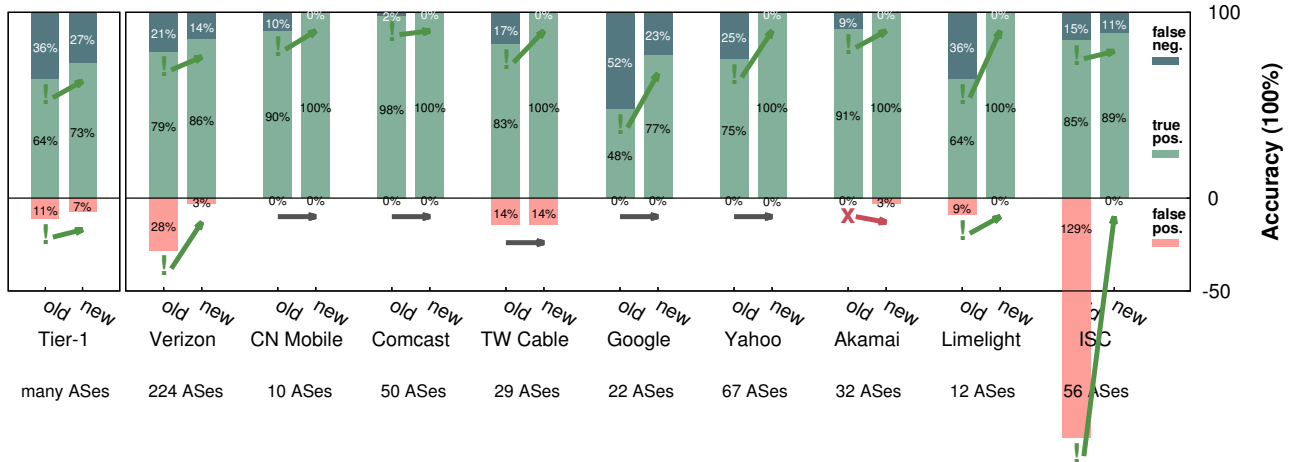


Figure 1: Comparison between our previous and current validation results.

category	orgs	ASes	addresses
total	36463	100%	49262 100%
multi-AS	4856	13%	17655 36%
single-AS	31607	87%	31607 64%
total	36463	49262	
routed	27802	100%	34472 100% 2.5B 100%
routing-complex	3165	11%	9835 29% 1.6B 64%
routing-simple	24637	89%	24637 71% 0.9B 36%
not routed	8661	14790	

Table 7: Organization distribution by number of ASes in total, and ASes in routing tables.

category	orgs	ASes	addresses
total	36463	100%	49262 100%
multi-AS	4856	13%	17655 36%
single-AS	31607	87%	31607 64%
total	36463	49262	
routed	27682	100%	34260 100% 2.5B 100%
routing-complex	3142	11%	9720 28% 1.6B 64%
routing-simple	24540	89%	24540 72% 0.9B 36%
not routed	8781	15002	

Table 8: Organization distribution by number of ASes in total, and ASes in routing tables (from the second RouteViews site).

AS organizations (second row). While most organizations use only a single AS (*single-AS*), about 36% of all ASes are assigned to multi-AS organizations. Since more than a third of ASes have other “sibling” ASes in the same organization, this finding suggests that an AS-to-organization map may be relevant to resolving AS-relationship discovery in routing [13].

However, some of the allocated ASes are not in active use and these “moribund” ASes may distort the picture. The above discussion is based on static information from WHOIS that includes moribund ASes. To evaluate ASes that are actually in use, we consider only the subset of ASes that are *routed*. We obtain routing tables from RouteViews [21] Oregon site; we see similar results from another vantage point located in Japan shown in Table 8 .

To account for moribund ASes, we define two terms, *routing-complex* and *routing-simple* organizations. After discarding moribund ASes, organizations that still have multiple ASes are called routing-complex, because they use multiple ASes to route; organizations that reduce to a single AS are called routing-simple. Note that routing-complex organizations must be multi-AS organizations, but not necessarily vice versa. Our main interest is in routing-complex organizations since they are the ones that actually use multiple ASes.

To account for the fact that different ASes have different degrees of *influence* on Internet, we approximate the influence of an AS by the number of addresses it announces, with the influence of an organization being defined as the sum of its AS’s influences. Note that the ideal way to define the influence of an AS is by the traffic it carries. Unfortunately, we know of no way to estimate the traffic carried by arbitrary ASes.

Table 7 (bottom portion) shows that *routing-complex* organizations, while only representing 11% of all routed organizations, have a big influence on the Internet, accounting for nearly two-thirds (1.6B or 64%) of all routed




Figure 2: Historical routability of Google ASes.

addresses. Specifically, routing-complex organizations that use more than 2 ASes make up only 4% of all routed organizations but announce more than half of all routed addresses, and routing-complex organizations with more than 5 ASes (some 1% of all routed organizations) announce about one third of all addresses. We see that these influential organizations have not just many addresses, but also many routed ASes.

5.2 Causes of multi-AS usage

Organizations use multiple ASes for many reasons. For our analysis, we group those into reasons that are *transient*, and those that are *persistent* and unlikely to go away. Mergers and acquisitions are examples of transient reasons, especially if we expect common business practices following a merger to include the consolidation of the ASes that were announced prior to the merger.

Persistent reasons are more varied, but usually result from some (internal or external) legal or policy pressure. An example of an internal policy decision is an ISP that chooses to use different ASes to implement internal routing policies; e.g., Verizon’s use of different ASes on different continents [36]. External policy constraints include cases where legal conditions of mergers require that certain business practices remain unchanged, or be handled independently post-merger. Regardless of the specifics, these kinds of policy constraints that are often in place for years, and so we label them “persistent”.

We next summarize our inference about both transient and persistent multi-AS usage of six organizations. To this end, we use n_p (where $p \in \{80, 100\}$) to denote the number of top ASes that announce $p\%$ of an organization’s addresses. Thus, n_{100} tells how many ASes are routed in total, while n_{80} focuses on the “core” ASes, where we typically see stable, policy-based ASes.

To illustrate multi-AS usage, Figure 2 gives a historical account of the routability of all ASes that are part of Google as of 2011-09-01 (see [7] for additional

examples). ASes are stacked by AS index (sorted by number of addresses currently announced and then by the first date routed), with horizontal bars indicating the periods when the ASes are routed (darker bars indicating membership in n_{80}). The first time an AS is routed is called n_{100} *birth* and the last one is called n_{100} *death*. Similarly, the first time of n_{80} membership is called n_{80} *promotion* and the last one is called n_{80} *demotion*. From the graph, we can see that two ASes have been announcing 80% of the addresses for one year: Google’s main AS (AS15169, AS index: 1) and a WiFi-specific AS (AS36492, AS index: 2), suggesting a stable routing policy.

In contrast, transient AS usage is often the result of acquisitions followed by AS consolidation. Continuing with the Google example in Figure 2, in late 2006, Google acquired Youtube (AS36561, AS Index: 16); as can be seen, the number of addresses announced by this AS gradually decreased, and the AS finally disappears from BGP by April 2011. This change suggests that, over time, Google consolidated this service into its core infrastructure.

To contrast with the Google example, we also observed a case where routing policy decisions promote AS diversification: ISC. Although only one AS announces most of ISC’s addresses, we notice that since 2003, ISC is using more and more ASes. Examining these new ASes, we see that each announces a single /24 address block. This policy is consistent with the choice to associate a unique AS with each physical anycast location [18] and ISC’s operation of the anycasted F-root DNS server. This example illustrates how policy decisions can result in an increasing number of ASes per organization and that this type of multi-AS usage is likely to last.

(Appendix F quantifies the persistence of multi-AS usage.)

6. APPLICATIONS OF AS-TO-ORG MAP

In this section we illustrate how accounting for multi-AS organizations impacts our understanding of a number of Internet topology-related features.

6.1 Edge Coverage of an AS vs. its ORG

We first consider the effects of an organization-level view of the Internet on our understanding of the network’s edge. Recall that routing-complex organizations that use multiple ASes control a majority of Internet edge in terms of routed addresses. In addition, they typically assign different ASes to specific geographic regions, mainly for implementing certain routing policies. Faced with the problem of inferring the geographic coverage of such a routing-complex organization, researchers must first identify the organizations’ ASes, then they need to obtain the addresses announced by




Figure 3: Edge Coverage at Org-level vs. AS-level for routing-complex organizations.

each AS, and finally, they have to geo-locate those addresses. Underestimation of an organization’s addresses, either due to equating the organization with its main AS or due to an incomplete AS-level view of that organization, will lead to an incomplete view on its geographic coverage.

To quantify this underestimation, Figure 3 shows the effects of having access to an AS-to-organization mapping on each of these steps. Figure 3(a) illustrates the incomplete view with respect to address coverage. For each routing-complex organization, we compare the number of addresses announced by the AS that announces the most addresses (i.e., main-AS view) to the number announced by all its ASes (i.e., org-level view). We calculate the percentage of addresses not announced by its main AS (*missed address coverage*), and show the cumulative distribution of organizations based on this percentage. The higher the missed address coverage, the more incomplete the view. As can be seen from Figure 3(a), the address coverage of almost all organizations is incomplete, missing between 1% to 91% of the addresses. More specifically, when considering an organization’s main AS only, then nearly one-third of all organizations (933) miss a significant portion (40% to 91%) of addresses.

To quantify the underestimation with respect to geographic coverage, we count the cities where these addresses are located. We identify the city of each address using the MaxMind’s CityLite geo-location database [19]. This dataset provides worldwide coverage, and claims that for the U.S., 79% of the addresses are mapped with an accuracy of less than 25 miles. With the help of this database, we identify the address locations for

2,631 routing-complex organizations (it does not have locations for 534 organizations, about 17%) and compute the *missed city coverage* in a similar fashion as the *missed address coverage* earlier. Figure 3(b) shows that nearly half of the mapped organizations (1,132) have an incomplete geographic coverage, missing at least one city. Importantly, when reducing an organization to its main AS, a fifth of the organizations (540) miss a significant portion of the cities (40% to 90%).

6.2 IXP Coverage of an AS vs. its ORG

To illustrate the value that an AS-to-organization map has on studying Internet IXP peering-related issues, we first show that the organizations that use multiple ASes to peer at IXPs (*peering-complex* organizations) have a big influence on the Internet. We then compare the main AS view and corresponding organization-level view of peering-complex organizations and discuss limitations of the former.

6.2.1 Importance of Peering-Complex Organizations

Large ISPs often use multiple ASes to implement routing policies for different geographical regions (Section 5.2). As a result, they tend to peer only with certain ASNs in certain locations. Similarly, large content and hosting providers also use different ASNs for, say, different continents. These ASes then peer with local access networks at close-by IXPs to save transit cost and optimize end user experience. Clearly, these are situations where the main AS view of an organization yields limited visibility into the true geographic reach of the corresponding organization and underestimates the geographic coverage resulting from an organization-level view.

To quantify the prevalence of multi-AS peering, we apply our AS-to-Org mapping to previously obtained AS-level IXP peering matrices [4]. These inferred peering matrices represent the current state-of-the-art but are known to be incomplete, and we will comment below on how this incompleteness may affect our observations. The 2009 dataset lists 2,840 ASes, and we map them to 2,503 organizations using our AS clustering method. Nearly two-thirds of these organizations are *routing-simple* and do not concern us here. In the following, we examine how the remaining 882 *routing-complex* organizations affect our view of IXP peering.

Figure 4(a) shows the cumulative distribution of organizations as a function of how many ASes they use to peer at IXPs. As can be seen, most (715 organizations of the 882) of these routing-complex organizations peer with only one AS (i.e., are *peering-simple*). Although only about 19% of these routing-complex organizations use multiple ASes to peer at IXPs (i.e., are *peering-complex*), Figure 4(b) shows that these peering-complex organizations have a large influence on the

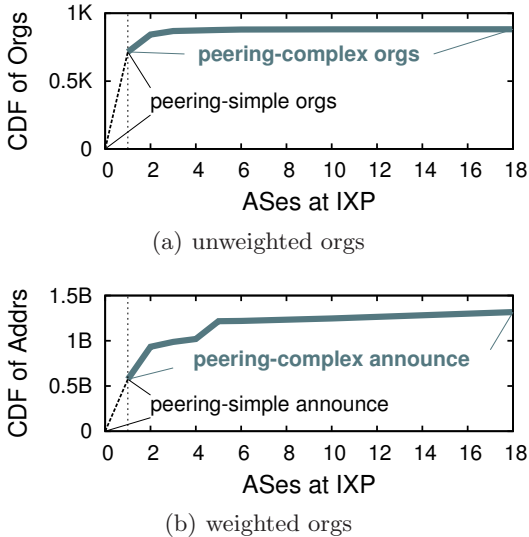


Figure 4: Cumulative distribution of unweighted/weighted peering-active organizations by number of ASes used to peer.

Internet. Approximating again influence by the number of announced addresses, Figure 4(b) shows that these peering-complex organizations account for more than half (0.74B or 56%) of all routed addresses (that were announced by both peering-complex and peering-simple organizations as of 2011-09-01). Note that the incompleteness of our IXP data suggests that our estimates are likely lower bounds on the number of peering-complex organizations.

Two examples illustrate peering-complex organizations as inferred from our 2009 data. The largest number of ASes used in peering is 18 by Comcast. These 18 ASes peer at 13 IXPs located in 12 cities, mostly in North America and Europe. The biggest organization, in terms of number of addresses, is China Telecom that uses 3 ASes to peer. These 3 ASes peer at 18 IXPs located in 16 cities in Europe and North America.

6.2.2 Implications of Peering-Complex Organizations

Next we quantify the degree of underestimation that results from reducing peering-complex organizations to the commonly-used main AS view. To this end, we focus on peering-related quantities such as geographic reach (i.e., number of IXPs or number of cities with an IXP), number of peers, and number of peering links.

For each organization, we measure the underestimation in each of the above three aspects by extracting subgraphs that represent both the organization-level view and the corresponding main AS view. Here, the notion of “main AS” is metric dependent, and refers to having the most IXPs, or peers, or links, respectively. An organization’s subgraph is the subset of the IXP

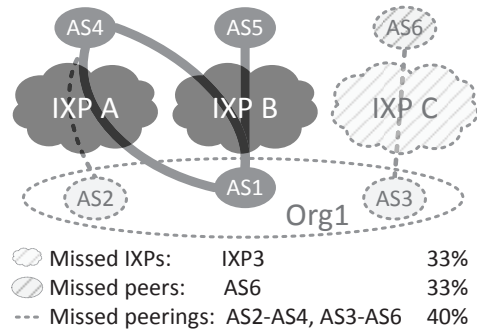


Figure 5: The different IXP peering view from the whole organization’s perspective and from the main AS’s perspective.

map that only contains the IXPs, peers and links related to it. The subgraph for the main AS consists of the IXPs, peers, or links associated with the main AS only. Note that the latter is always a subset of the organization’s subgraph.

Figure 5 gives an example of these measures. *Org1*’s subgraph is shown as the whole figure, while if AS1 is its main AS, then the gray portions refer to its main AS subgraph. *Org1*, consisting of three ASes, peers at three IXPs, with three ASes and five peering links in total. In contrast, the main AS, AS1, peers only at two IXPs, with two ASes and three peering links.

We quantify the degree of underestimation by counting the fraction of IXPs, peering ASes and peering links, resp. that one will miss by considering only the main AS’s graph. For instance, in the case of the example in Figure 5, we can see that the main AS view misses 33% of the IXPs, 33% of the peers and 40% of the peerings.

Figure 6 gives for the different metrics the cumulative distribution of all 167 *peering-complex* organizations by their degree of underestimation. For example, in terms of (geographic) coverage, we observe that the main-AS view often provides a limited perspective of the (geographic) coverage of the corresponding organization-level view—about one-third of all organizations are missing 20% of IXPs and cities (see Figure 6(a) and Figure 6(b)). Turning to peers, Figure 6(c) shows that the majority (144) of organizations miss some peers and about one-third of organizations miss more than 20% of the peers. These omissions can lead to a false inference of an organization’s peering strategy. Lastly, the underestimation is worst with respect to peering links. Figure 6(d) shows that almost all organizations are missing some peering links when compared to the main AS view, with half of the organizations missing more than 20%. The underestimation of peering links can result in underestimates of the organizations’ IXP-specific connectivity fabrics.




Figure 6: IXP Peering at Org-level vs. AS-level for peering-complex organizations.

7. CONCLUSIONS

This paper describes a new algorithm that automatically yields AS-to-organization maps of the Internet. We show these maps improve accuracy compared to current maps, a result stemming from the use of company subsidiary data contained in the annual U.S. SEC Form 10-K filings,

and better clustering methods, each steps that link or separate previously incorrect clusters. We validate our new AS-to-organization map against a “best-effort” ground truth. Finally, we show that accounting for all of an organizations ASes provides much better representation of organization properties such as size, geographic footprint, and IXP peerings, compared to the commonly-applied “main AS” views.

Our work provides a basis for gaining a deeper understanding of the business relationships that exist in today’s Internet between the ASes of an organization and between organizations. We expect these relationships to continue to evolve with the Internet, particularly as businesses become more heterogeneous and interconnected (as ISPs distribute content, and content providers deploy networks), increasing the relevance of approaches that reveal these business relationships in the Internet’s main players. We believe that these more complete, organizational views will enable a more holistic, and ideally more realistic, study of the impact of real-world threats to the Internet, be they physical [3], economic [29], or political [20].

8. REFERENCES

- [1] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance in complex networks. *Nature*, July 27 2000.
- [2] ARIN. Introduction to ARIN’s database. <https://www.arin.net/knowledge/database.html>, March 2010.
- [3] Associated Press. Fire in Baltimore snarls Internet traffic, too. *New York Times*, July 20 2001.
- [4] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. IXPs: mapped? In *IMC’09*, New York, NY, USA, 2009. ACM.
- [5] Time Warner Cable. Locations. <http://www.timewarnercable.com/corporate/about/careers/locations.html>, November 2011.
- [6] Xue Cai, John Heidemann, Balachander Krishnamurthy, and Walter Willinger. Towards an AS-to-Organization map. In *IMC’10*, Melbourne, Australia, November 2010. ACM.
- [7] Xue Cai, John Heidemann, Balachander Krishnamurthy, and Walter Willinger. Accurate AS-to-organization mapping and its implications (extended). Technical Report ISI-TR-679, USC/Information Sciences Institute, 2012. <ftp://ftp.isi.edu/isi-pubs/tr-679.pdf>.
- [8] Amit Chandel, Oktie Hassanzadeh, Nick Koudas, Mohammad Sadoghi, and Divesh Srivastava. Benchmarking declarative approximate selection predicates. In *ACM SIGMOD’07*, New York, NY, USA, 2007.
- [9] James H. Cowie, Andy T. Ogielski, BJ Premore, Eric A. Smith, and Todd Underwood. Impact of the 2003 Blackouts on Internet Communications. https://www.renesys.com/tech/presentations/pdf/Renesys_BlackoutReport.pdf, 2003.
- [10] Danny Dolev, Sugih Jamin, Osnat Mokryn, and Yuval Shavitt. Internet resiliency to attacks and failures under BGP policy routing. *Computer Networks*, 50(16):3183–3196, 2006.
- [11] Jeff Ehling. Time-warner cable leaving houston. <http://abclocal.go.com/ktrk/story?section=news/local&id=4423119>, August 2006.
- [12] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *Proc. of ACM SIGCOMM*, Cambridge, MA, USA, September 1999. ACM.
- [13] Lixin Gao. On inferring autonomous system relationships in the Internet. *ACM/IEEE Transactions on Networking*, December 2001.
- [14] Ramesh Govindan and Anoop Reddy. An Analysis of Internet Inter-Domain Topology and Route Stability. In *in Proc. IEEE INFOCOM*, pages 850–857, 1997.
- [15] Donald Greenlees and Wayne Arnold. Asia scrambles to restore communications after quake. http://www.nytimes.com/2006/12/28/business/worldbusiness/28iht-connect.4042439.html?_r=1, December 2006.
- [16] J. Hawkinson and T. Bates. Guidelines for creation, selection, and registration of an

- autonomous system (AS). RFC 1930, Internet Request For Comments, March 1996.
- [17] Young Hyun, Andre Broido, and k. c. claffy. Traceroute and BGP AS path incongruities. Technical report, UCSD CAIDA, 2003.
- [18] ISC. Origin ASN for anycasted services. <http://www.isc.org/community/blog/201109/origin-asn-anycasted-services>, September 2011.
- [19] MaxMind. Geolite city. <http://www.maxmind.com/app/geolitecity>, March 2012.
- [20] Danny McPherson. Internet routing insecurity: Pakistan nukes YouTube? Security to the Core: the Arbor Networks Security Blog, February 2008.
- [21] David Meyer. University of Oregon Route Views project. <http://www.routeviews.org>.
- [22] M. Newman, A.-L. Barabási, and D.J. Watts. *The Structure and Dynamics of Networks*. Princeton University Press, 2006.
- [23] M. Omer, R. Nilchiani, and A. Mostashari. Measuring the resilience of the global Internet infrastructure system. In *Proc. of the 2009 IEEE International Systems Conference*, 2009.
- [24] Packet Clearing House. PCH INOC-DBA. <https://www.pch.net/inoc-dba/>, April 2010.
- [25] Regional Internet Registry. <http://www.afrinic.net/>, <http://www.apnic.net/>, <http://www.arin.net/>, <http://www.lacnic.net/>, <http://www.ripe.net/>, November 2009.
- [26] Y. Rekhter and T. Li. A border gateway protocol 4 (BGP-4). RFC 1771, Internet Request For Comments, March 1995.
- [27] Renesys. Internet-Wide Catastrophe—Last Year. http://www.renesys.com/blog/2005/12/internetwide_nearcatastrophe1a.shtml, December 2005.
- [28] Renesys. Con-Ed Steals the 'Net. http://www.renesys.com/blog/2006/01/coned_steals_the_net.shtml, January 2006.
- [29] Renesys. Sprint and Cogent Peer. <http://www.renesys.com/blog/2006/11/sprint-and-cogent-peer.shtml>, November 2006.
- [30] RIPE. RIPE database query reference manual. <http://www.ripe.net/db/support/query-reference-manual.pdf>, November 2009.
- [31] Stephen E. Robertson, Steve Walker, Micheline Hancock-Beaulieu, Mike Gatford, and A. Payne. Okapi at TREC-4. In *TREC'95*, 1995.
- [32] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2003.
- [33] Gerard Salton and Christopher Buckley. Term-weighting approaches in automatic text retrieval. In *Information Processing and Management*, pages 513–523, 1988.
- [34] SEC. Researching public companies through EDGAR: A guide for investors. <http://www.sec.gov/investor/pubs/edgarguide.htm>.
- [35] The National Academic Press. The Internet under crisis conditions: Learning from the impact of September 11. <http://books.nap.edu/openbook.php?isbn=0309087023>, 2003.
- [36] Verizon. Verizon business policy for settlement-free interconnection with internet networks. <http://www.verizonbusiness.com/terms/peering/>, September 2011.
- [37] H. M. Walker. A theory for record linkage. *Journal of the American Statistical Association*, 1969.
- [38] WIRED. Oops! AT&T Blackhole Was 4Chan's Fault. <http://www.wired.com/threatlevel/2009/07/4chan/>, July 2009.
- [39] J. Wu, Y. Zhang, Z.M. Morley, and K.G. Shin. Internet routing resilience to failures: Analysis and implications. In *Proc. ACM CoNEXT'07*, New York, US, 2007.

APPENDIX

A. DATASET DIFFERENCE BETWEEN MAIN PAPER AND APPENDIX

The dataset used in Appendix C, E and F is different from the one used in the main body of this paper. The two main differences are as follows. First, in the old data used in this appendix, RIPE WHOIS data is acquired in bulk directly from the provider and is not fetched from the RIPE website. This bulk dataset misses most of the details about contact information, including phone numbers and e-mail addresses—these are the very details we use to establish clustering in our approach.

Second, in terms of the input data used for this appendix, 10-K links are not manually verified and pruned, and include all 8,706 organizations; that is, not just the manually verified 50 organizations described in Section 3.2.3. We use *all10K* to annotate this difference. For example, input $\mathcal{I}_{attr+all10K}$ is different from $\mathcal{I}_{attr+10K}$ in the sense that the former contains all 10-K links, while the latter only contains verified and pruned ones.

Effect on training (Appendix C) Our use of this old dataset for training could potentially result in poor weights. However, we claim that our training is reasonable in spite of this risk. For one, although training is done with this bulk RIPE dataset, the WHOIS datasets from the other four RIRs are the same in both the current and old input. Moreover, the other four RIRs provide full data, supporting discovery of reasonable weights.

The old 10-K dataset also causes the inputs with 10-K to underperform the ones without it (see Table 9), and it was the motivation for our decision to improve the use of 10-K information through a largely manual process.

This difference between the two datasets means that one should not compare results with and without 10-K in Table 9; the 10-K data used there is not as accurate as our downselected and manually verified 10-K data.

Our primary goal here is to provide enough details about our training *process* to conclude that it is thorough enough for the purpose at hand. Although the use of old data here may alter the quantitative conclusions, it does not affect the description of the training process.

Effect on validation results (Appendix C) The two changes to the old input data were done to improve false-negative rates, so the newer data should produce better results when compared to PCH. Thus we argue that Table 10 is therefore a lower bound on the accuracy of our approach.

Effect on multi-AS usage (Appendix F) Because the new 10-K data only covers 50 organizations, we do not expect it will affect our classification results of multi-AS usage (Figure 17) much. The additional data in RIPE may change the shape of Figure 17, causing previously single-AS or small, multi-AS organizations to be grouped into bigger, multi-AS organizations. Thus our classification results should be considered as a preliminary result and not a final conclusion.

B. ATTRIBUTE EXTRACTION AND STANDARDIZATION IN DETAIL

In Section 3.1.2, we summarized how we selected and normalized the AS attributes used for clustering. This section gives more details, highlighting what aspects of the process differ from our prior work [6].

Extract Raw Attributes: We extract raw attributes from *WHOIS*, following chains of AS- to org- to contact-record as necessary.

Canonicalize to Simple Attributes: Unlike OrgID and contact IDs, phone and email attributes often contain details that make similar records appear dissimilar. For example, telephone numbers need to be amended with country code and stripped of extensions. To canonicalize email records, we discard the user portion and keep only the distinguishing, right-most part of the domain address. We identify that portion with a manually-built list of more than 6k suffixes using longest-suffix matching.

Discard Generic Attributes: A number of attribute values are *generic*, shared by unrelated ASes. Examples of generic attributes are public email services like Gmail and Hotmail. Used blindly, generic attributes will link unrelated organizations into large, incorrect clusters. In addition to public e-mail providers, we identify eight generic OrgID attributes (for RIRs, IANA, and NICs) and 120 contact IDs (for these, plus tens of outsourcing companies and a few ISPs that man-

age customer networks). In total, we discard 179 phone numbers and 141 e-mail domains.

C. TRAINING IN DETAIL

In Section 3.1.3, we briefly described how we tuned the weights for the different attribute types. In this section, we give a more detailed description, describing our training method in Section C.1, and providing details about our training results in Section C.2.

C.1 Training Methodology

We first list the parameters to optimize (Section C.1.1), then set aside a training set of about 10,000 ASes (Section C.1.2), and finally define the objective function that we attempt to optimize (Section C.1.3) by relying on a parallel hill climbing algorithm (Section C.1.4).

C.1.1 Parameters

We consider several parameters that greatly affect our AS-to-organization mapping results and list them below. Our focus here is to determine the best possible weight vector \hat{w} for these parameters to improve our results.

First, how specific or general should attributes be? We group attributes into either the *4attr* or *66attr* set.

Second, what data sources should we use? Besides *WHOIS* data, we can choose either to use or omit the 10-K data.

Third, when should we merge clusters or leave them distinct? We must determine a *cutting threshold* ϵ for hierarchical clustering; this threshold is used to decide when a similarity score is too low to keep two clusters together. As stated in Section 3.1.5, our similarity score is proportional to sum of the weights. Thus, we first define two fixed cutting thresholds, a conservative one with $\epsilon = 0.01$ and a more liberal one with $\epsilon = 0.001$, and then let the training algorithm walk through different weight vectors. After the best weight vector is selected, we normalize it and adjust the cutting threshold accordingly.

We consider each of the combinations of the four attribute sets (*4attr*, *66attr*, *4attr+all10K*, *66attr+all10K*) and two cutting thresholds ($\epsilon = 0.01$ and $\epsilon = 0.001$) and optimize weights accordingly.

C.1.2 Sample Dataset

We first need to select a training sample of ASes. An ideal training set should be verifiable, representative and computable. *Verifiable* means that we have sound ground truth to evaluate the clustering results in order to guide training. *Representative* means that the sample contains an appropriate subset of all ASes, so both false positives and false negatives can be captured. *Computable* means we can evaluate a training run reasonably quickly on a commodity computer. Finding

that the memory requirement for clustering represents the bottleneck for training, we adjust the size of the training set accordingly.

However, there are two problems with obtain the ideal training set. First, we have high-confidence ground truth only for ten organizations (the Tier-1 ISP and the nine organizations in Section 4.1). Thus, only results for these ten organizations are verifiable with respect to both false-positives and false-negatives. Second, clustering is very memory- and CPU-intensive if the dataset is large (memory $\propto N^2$, time $\propto N^3$ for N ASes in the training dataset). Clustering with some 50K ASes requires about 24GB memory and takes approximately 3 days on a large-memory computer. Due to limited access to this hardware and timing limits imposed by the large parameter space, we train only on a subset of the data.

To this end, we created a training sample with 9710 ASes, selecting about one fifth of the whole AS population. This size is small enough to make the analysis tractable—one clustering based on one weight vector requires about 1 GB of memory and takes about 20 minutes. One round of hill climbing walks through approximately 60 weight vectors before it converges. We run 32 rounds for each of the eight combinations of attribute sets and cutting thresholds. Thus in total, the training takes 20 minutes per case, with $60 \times 32 \times 8 = 15\text{k}$ cases, or 213 days of compute time. We carry this work out in parallel on 32 processors over about 7 days. To make the training sample verifiable, we begin by seeding it with all ASes known to be in the ten organizations (736 ASes in total). To make it representative, we then add about 9K additional ASes of “noise” as described below. We then train on this dataset to choose the best clustering scheme and parameters.

We select our parameters based on training with a purposefully-chosen subset of ASes. We considered two approaches to choose “noise” ASes to fill out the training set: select them randomly from all ASes, or select them preferentially in the sense of being close to, but not in, the ten organizations. We say an AS is *close* to the ten organizations if this AS is likely to be clustered with them under a general attribute set and a very low cutting threshold. In the case of preferential selection, we select all ASes that cluster with the ten organizations under the *4attr+all10K* attribute set with *0.0001* cutting threshold. This preferential approach biases the noise to make training more difficult, mainly because it is easy to accidentally cluster nearby ASes with known organizations. In Section C.2.1 we confirm that our biased noise produces more accurate parameters compared to using a purely random selection.

C.1.3 Objective Function

Here we define the *objective function* $f(\hat{w})$ used in

Section C.1.4 to judge what value for \hat{w} best reflects clustering quality. A simple function would sum false-positives and false-negatives for the ten organizations, but we found this approach to be very sensitive to outliers. To avoid this problem, we instead sum the quartiles of false-positives and false-negatives. Let fp_{Q_1} , fp_{Q_2} , and fp_{Q_3} be the first, second and third quartiles of false-positive rates validated by the ten organizations, and fn_{Q_1} , fn_{Q_2} , and fn_{Q_3} be the false-negative rate quartiles (see Equation (2) in Section 4.2 for false-positive/negative rate definition). The objective function is then defined as

$$f(\hat{w}) = \sum_{i=1}^3 fp_{Q_i} + fn_{Q_i}$$

and leverages the false-positive and false-negative rates for comprehensiveness of our objective and also discards outliers by using representative quartiles. Based on this definition, a lower $f(\hat{w})$ means better \hat{w} , and thus the goal of the hill climbing algorithm is to *minimize* $f(\hat{w})$.

C.1.4 Algorithm: Parallel Hill Climbing

The algorithm we chose is *parallel hill climbing*. The basic hill climbing algorithm starts from a random \hat{w} -value, iteratively tries to find a better one by changing one element of it and judging if the new one produces a better clustering result. The algorithm iterates until it reaches a local optimum where no improvements can be found around the final \hat{w} , even after an exhaustively search of nearby configurations.

Basic hill climbing is fast at finding a *local optimum*, but it cannot determine if that point is a *global optimum*. To increase the chances of finding a global optimum, we use *parallel hill climbing* which starts from multiple random \hat{w} -values. Provided it iterates long enough, parallel hill climbing finds local optima with high probability. With enough initial positions, parallel hill climbing will find a good global value with high probability, provided the parameter space is relatively smooth. As illustrated in Figure 8, this assumption appears to hold in our case.

We run 32 parallel hill climbing processes for each of the eight combinations of attribute set and cutting threshold. Table 9 shows the number of different \hat{w} -values examined in the training space for each combination in parenthesis. Each initial value searches a path through space until reaching a local optimum. With 32 rounds of searching, parallel hill climbing typically explore 1 to 2k weight vectors. Although our random walks cover less than 1% of the training space, as Figure 7 shows, the best (lowest) scores for all eight attribute/threshold combinations converge fairly quickly, usually after 9 rounds. Thus, we argue that 32 rounds are sufficient to find a local optimum and to determine a set of “good” parameters.




Figure 7: Converging training results with parallel hill climbing.

Input	Cutting Threshold	
	0.01	0.001
4attr	44.5 (1.5k)	45.0 (2.2k)
4attr+all10K	45.5 (1.1k)	47.0 (1.0k)
66attr	95.0 (2.5k)	75.5 (2.5k)
66attr+all10K	65.5 (2.2k)	63.0 (2.2k)

Table 9: Summary of training results. Best score: 44.5. Numbers of weight vectors examined are in parenthesis.

C.2 Details of Training Results

In this section, we first present the best parameters we found and then discuss how different parameters affect the results.

C.2.1 Best parameters

Table 9 gives a summary of our training results. The best attribute set/threshold combination is $4attr+0.01$ with score 44.5; the combination $4attr+0.001$ has a very similar score, indicating that the result is not very sensitive to the value of the cutting threshold. The corresponding weight vector is given by $\hat{w}=\{3, 0.4, 0.4, 0.2\}$, and if we normalize so the sum of the weights is 1, then $\hat{w}=\{0.75, 0.1, 0.1, 0.05\}$, with $\epsilon = 0.0025$. Either way, this solution quantifies the importance of the different attribute types. It emphasizes the importance of OrgID, and downplays contact ID, phone and email. This result confirms our expectation that while OrgID is very informative for the purpose of clustering, email is much less useful. Since OrgIDs are intended for common administrative management, they are unlikely to cause false-positives; on the other hand, since contact information can be registered by outsourcing third parties, they can easily introduce false-positives.

To verify that our method of training dataset selection is appropriate, we also trained using a purely random selection of 9K noise ASes. While the best

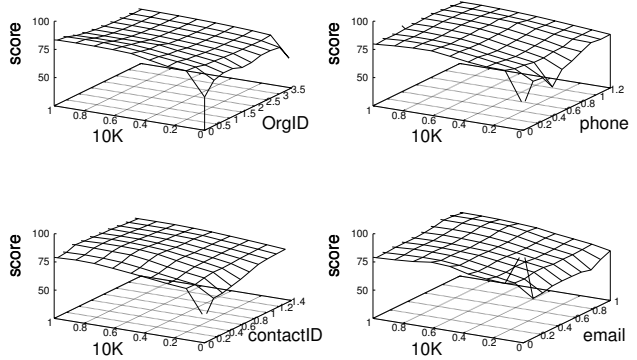


Figure 8: Parallel Hill Climbing with attribute set $4attr+all10K$ and cutting threshold 0.01 . company subsidiary information (10k) is always shown on the y axis.

score resulting from training with this input data is better than that resulting from training with preferential noise (34.5 instead of 44.5), the resulting weights ($\hat{w}_R = \{0.4, 0.28, 0.24, 0.08\}$) and threshold ($\epsilon = 0.0004$) perform much worse when applied to the whole dataset, scoring 86.5 compared to 49.5 with preferentially-derived parameters. This result confirms that our biased, more challenging training dataset improves overall accuracy by finding more effective weights.

C.2.2 Parameter discussion

In this section, we discuss several factors affecting our training results and understand why other parameters are not preferred.

Attribute generalization Attribute/threshold combinations with 66 attribute types perform much worse than the ones with four attribute types. This is because dividing a general attribute type into many specific sub-types may break clustering links and thus may lead to higher false-negative rates. With the 4attr (or 4attr+all10K) set, *administrative* contact email @example.com (belongs to AS1) and *technical* contact email @example.com (belongs to AS2) are of the same type (both belong to *email* type), thus they will be compared with each other. Since these two emails are of the same value, AS1 and AS2 will be linked together. However, with the 66attr (or 66attr+all10K) set, *administrative* email and *technical* email are of different types. As mentioned in Section 3.1.4, different types of attributes are orthogonal, thus they will not be compared and AS1 and AS2 will not be linked.

Attribute weights Although we chose a best value for the weight vector \hat{w} in Section C.2.1, we see the training space is fairly flat. Figure 8 visualizes the training space for attribute set $4attr+all10K$. As can be seen, the score

is not very sensitive to the weight changes. Instead, it is more sensitive to the selection of different number of attribute types as shown in Table 9.

D. MANUAL VERIFICATION AND PRUNING IN DETAIL

In Section 3.2.3 we briefly discussed how we manually verified and used 10-K links for some 50 purposefully selected organizations. In this section, we present more details about this process.

We select 50 organizations (about 0.6% of the 8,706 organizations) intentionally to favor those that are relevant in the real world and are important to the Internet’s ecosystem. In particular, we select 38 large, computer-related organizations from the 2011 Fortune 500 list and add 12 large ISPs that are not included in that list. In terms of Fortune 500 companies, we included all organizations (38 in total) in the following six Internet-related industries: telecommunications (e.g., Verizon, Sprint), Internet services (Amazon, Google), computers (Hewlett-Packard, Apple), software (Microsoft, Oracle), IT services (IBM, Computer Sciences Corporation), and communication equipment (Cisco). We then added 12 organizations that are not on the Fortune 500 list but are important players in the Internet, including large Tier-1 and Tier-2 ISPs such as Level 3 and Cogent.

The complete list of these 50 organizations is:

1. Telecommunication (14 companies): AT&T, Cablevision, Charter Communications, Comcast, DirecTV, DISH Network, Liberty Global, NII Holdings and Telephone & Data Systems, Qwest, Sprint, Time Warner Cable, Verizon, and Virgin Media,.
2. Internet service (5 companies): Amazon, eBay, Google, Liberty Media, and Yahoo.
3. Computer (5 companies): Apple Inc., Dell Inc., Hewlett-Packard, Pitney Bowes, and Xerox.
4. Software (3 companies): Microsoft, Oracle, and Symantec.
5. IT service (5 companies): AimNet Solutions, Cognizant, Computer Sciences Corporation, IBM, and SAIC Inc.
6. Communication equipment (6 companies): Avaya, Cisco, Corning Inc., Harris Corporation, Motorola, and Qualcomm.
7. Other (12 companies): Akamai, Citigroup, Cogent, Equinix, Gannett, Internap, Limelight, Savvis, SunGard, VeriSign, Vonage, and XO Communications.

Of the 1817 links that the automated clustering produced for these 50 organizations, we verified and kept 1226 links, dropping 591. To verify the correctness of a link, we manually compared the AS name with the subsidiary name, using additional information from

WHOIS and public web pages where available. For example, we verified and kept the link between AS36561 (*YouTube, Inc.*) and Google’s subsidiary (*YouTube, LLC*), because AS36561 registered with the same address as Google, and has contacts with e-mail domain *google.com* and *youtube.com*. In contrast, we eliminated the link between AS4616 (*Information Technology Services*) and Google’s subsidiary (*Google Information Technology Services LLC*), because AS4616 actually belongs to Hong Kong Polytechnic University according to its WHOIS record.

E. VALIDATION WITH BROADER COVERAGE (PCH)

In this section, we use PCH’s manually generated AS-to-organization map (alluded to in Section 2) to validate our work with a dataset that provides broader coverage than our carefully chosen validation datasets. We evaluate this related work and demonstrate its incompleteness (Section E.2). Although incomplete, we use it to test our clustering algorithms for false-negatives (Section E.3), and present the results in Section E.4.

E.1 Validation Dataset

The PCH dataset (referred as T_{pch}) is a database that relies on *voluntary* contributions from network operations personnel in many different organizations and is maintained by PCH to facilitate communication among the different players interested in a smooth functioning of the Internet. Compared to our other validation datasets, T_{pch} covers many more organizations (960 in PCH, compared to a total of 110 organizations in our other datasets). Especially, T_{pch} is more diverse in terms of organization sampling (similar as $T_{randall}$), mainly because it covers many “small” organizations with fewer ASes (mean cluster size is only 2). We also expect it to be more unbiased than, for example, $T_{randtop}$.

The PCH data is a table with three columns: AS, *shortorg* and *longorg*. *Longorg* and *shortorg* are full and abbreviated names of the organization to which the AS is assigned; e.g., “Internet Systems Consortium, Inc.” and “ISC”. There is no strict format for *shortorg* and *longorg*, and not every AS has both *shortorg* and *longorg*. *Longorgs* are 20 times more frequent than *shortorgs*, however they are usually verbose and contain details that make string matching hard.

Because *longorgs* make clustering difficult, we identify AS clusters in T_{pch} by *shortorg*. ASes with the same *shortorg* are clustered together and identified as belonging to one and the same organization. As Table 2 shows, this process results in about 2K ASes grouped into 960 clusters.

Although PCH data covers many more organizations than our other validation sets, it covers fewer ASes for

each organization, and these ASes may fall into different shortorg clusters (see Section E.2). This drawback of PCH data poses challenges for our validation efforts. In particular, we cannot validate false-positives because the ground truth itself is incomplete.

E.2 Evaluation of PCH Dataset with Strong Ground Truth

Because the PCH dataset is the result of a largely voluntary effort, we expect that it will be less complete than the datasets we have built ourselves. Therefore, we first evaluate the completeness of the PCH dataset before using it to judge the accuracy of our clustering algorithm.

To assess the quality of PCH dataset, we compare it with T_{tier1} and T_{9org} . We use three different validation metrics to achieve three different objectives: gain some intuition, obtain error bounds, and calculate correction factors.

We first use the same validation metric described in Section 4.2 (Figure 9a). This definition is the most intuitive one, and it allows us to compare PCH’s data quality with our results. For any given ground truth AS cluster C_0 in T_{tier1} or T_{9org} , we select the cluster C_1 in T_{pch} that has the largest overlap with C_0 , and then compare them (Figure 9a). However, this method only gives us an approximate assessment of PCH’s data quality. For example, it ignores other clusters that overlap with C_0 , thus misses both true-positive and false-positive ASes in these clusters. We therefore use a second definition to obtain some bounds on the errors.

The second approach considers not just the biggest cluster, but selects *all* overlapping clusters (say, C_1 and C_2 in Figure 9 b). We then take the union of all these overlapping clusters and compare the resulting cluster with the ground truth cluster. Given that this approach covers all clusters, it provides a lower bound for missing ASes (false-negatives) and an upper bound for wrong assertions (false-positives).

Later in Section E.3, we add a third definition to evaluate the PCH data. The purpose of this evaluation is not to assess the quality of PCH data, but rather to obtain correction factors that can be applied when we validate our results using PCH data as ground truth. The need for this third metric arises from the incompleteness of the PCH data.

Figure 10 shows the evaluation results for the PCH data. Using the *biggest* and *all* metrics, we observe few false-positives, but many false-negatives. Checking the results for the *biggest* metric, only the Tier-1 ISP has a small false-positive rate, with all other organizations, except for Comcast and Time Warner Cable, missing more than 50% of the ASes. When comparing these results with the ones in Figure 1, we see that our clustering approach outperforms the clustering pro-




Figure 11: The adjusted false-negative rate.

vided by PCH.

The second definition, the (*all*) metric, produces an upper bound for false-positives and a lower bound for false-negatives. We see that except for Time Warner Cable, all organizations have at most 3% false-positives. The Time Warner Cable situation is caused by a single AS with incorrect information in PCH that introduces 15 false positives. As for a lower bound on false-negatives, 6 organizations miss at least 90% of the ASes; this confirms the incompleteness of PCH data.

We thus conclude that T_{pch} is relatively correct but incomplete. For the purpose of validating our clustering results, we find that T_{pch} is suitable for assessing false-negatives; that is, if two ASes are in the same cluster in T_{pch} , then they should be in the same cluster in our results.

E.3 Validation of Our Results with PCH

To validate our results with T_{pch} , we introduce a new definition of false-negative rate fn_p to address PCH’s incompleteness. Note that since the PCH data is incomplete, we do not validate the false-positive rate with T_{pch} for our results.

The challenge that the PCH data poses is how to relate AS clusters to organizations. In the cases of T_{tier1} and T_{9org} , each AS cluster is associated with exactly one organization. However, due to PCH’s incompleteness, one organization can have multiple AS clusters. Which cluster should we choose to compare with the one in our result? Furthermore, how do we identify multiple AS clusters that belong to the same organization in the first place?

This challenge makes it hard to calculate errors for each organization as we did with T_{tier1} and T_{9org} . Instead, we compute a single false-negative rate for all clusters. To compute this single rate, we introduce the *pair* metric that counts AS pair links rather than ASes (see Figure 9 c)). In contrast with the previous two metrics, only this new metric can capture the clustering results when aggregated into a single rate.

To formally define this single false-negative rate, consider all AS pairs, x and y , and set

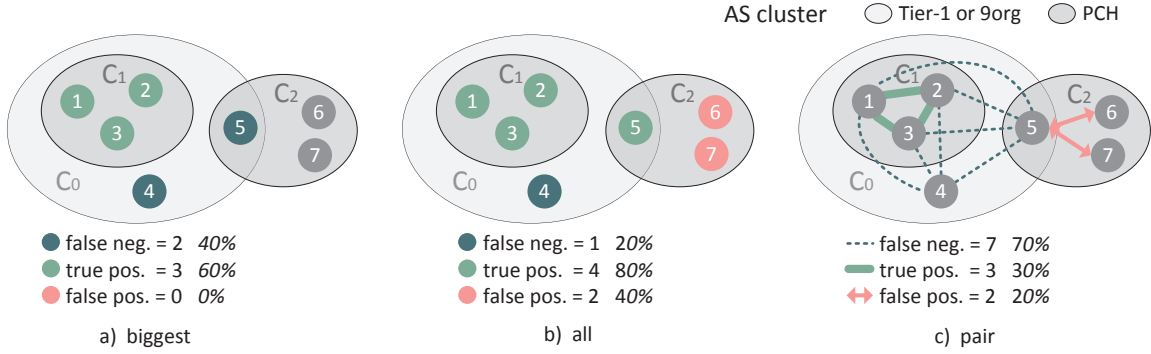


Figure 9: Three definition of validation metrics.




Figure 10: Evaluation of PCH dataset, compared with the Tier-1 ISP and 9 organizations.

$$\begin{aligned}
 P_{pch} &:= \text{ASes}(x, y) \text{ such that} \\
 & \text{org}_{pch}(x) = \text{org}_{pch}(y), \\
 & \text{org}_{pch}(x) \in T_{pch}
 \end{aligned}$$

and

$$\begin{aligned}
 P_{ours} &:= \text{ASes}(x, y) \text{ such that} \\
 & \text{org}_{ours}(x) = \text{org}_{ours}(y), \\
 & \text{org}_{ours}(x) \in R_{ours}
 \end{aligned}$$

where $\text{org}_{pch}(x)$ and $\text{org}_{ours}(x)$ are the corresponding organizations of AS x in PCH and in our results, respectively. We define the *relative* false-negative rate of our results compared with the PCH data as

$$fn_p = 1 - \frac{|P_{pch} \cap P_{ours}|}{|P_{pch}|}$$

This relative rate fn_p only tells us how well we did compared with the PCH data (but not compared to the ideal ground truth). To account for this effect, we introduce two correction factors k and b to obtain an approximate estimation of the *absolute* false-negative rate. Figure 11 illustrates the correction process. We first obtain the relative false-negative rate (fn_p) by com-

	P_{tier1}	P_{ideal}	P_{9org}
fn_p		29%	
fn_p^*	64%*		44%*
k	30%		15%
b	15%		45%
$1 - k'$	67%		41%

Table 10: Validation results by PCH

paring our results with PCH data. We then compute (i) how much of the “ideal ground truth” PCH covers (k), and (ii) how much of the ideal ground truth PCH misses but our results cover (b). Finally we calculate the corrected false-negative rate (fn_p') using the equation shown in the graph. The corrected false-negative rate takes PCH’s incompleteness into consideration by re-weighting the relative rate (multiplication by k) and amending the missing portion (addition of b).

However, since we do not have the “ideal ground truth”, we approximate it using either P_{tier1} or P_{9org} . We caution that this approximation may introduce errors.

E.4 Validation Results

Table 10 shows the relative and the corrected false-negative rates of our results for the PCH data. We missed 29% of the AS pairs in PCH. If corrected by either P_{tier1} or P_{9org} , the false-negative rates increase to 64% or 44%. Note that $1 - k'$ is actually the relative false-negative rate of our results compared with the ideal ground truth (see Figure 11). Since we use P_{tier1} or P_{9org} to approximate the ideal ground truth, $1 - k'$ shows how many AS pairs we missed for our 10 organizations. We see that the corrected false-negative rates are consistent with the validation results for either T_{tier1} or T_{9org} .

In summary, we conclude that our clustering results are consistently more accurate when compared to a much broader set of ground truth.

F. PERSISTENCE OF MULTI-AS USAGE

In Section 5.2, we briefly examined why organizations use multiple ASes, classifying causes into being either transient or persistent in nature. In this section we look at these causes in more detail. We first develop a method to classify organizations to understand transient and persistent AS use (Section F.1). We then give two examples to illustrate our classification method (Section F.2) and subsequently verify the practicality and correctness of our method (Section F.3). Lastly, we apply our classification to all organizations identified by our AS-to-org mapping effort and demonstrate the *prevalence* and *persistence* of multi-AS usage (Section F.4).

F.1 Evolution of multi-AS usage

While our analysis of the current Internet ecosystem shows that many ASes are part of multi-AS organizations, we wish to understand if this finding is an artifact of today’s Internet, or if multi-AS usage is growing or shrinking. To this end, we examine the per-organization changes of ASes active in routing tables over time.

We evaluate the importance of an AS by counting the number of addresses it originates. A persistent AS will be the origin of prefixes for years, while a transient AS’s announcements will eventually vanish. We measure how many important ASes an organization has by counting the number of ASes that announce 100% and 80% of the addresses for each organization (denoted by n_p , where $p \in \{100, 80\}$), and examine the trend of this number. A constant number indicates persistent usage, while a changing number may indicate transient usage, except for possible “noise” that first needs to be quantified. Based on observed trends in these numbers, we classify organizations into four categories: inconsistent, constant, consolidating, diversifying. We also bound the number of organizations that will keep using multiple ASes. The method takes the following steps.

First, to establish a base to compare with historical

snapshots, we obtain a current address set A for each organization. We match all IPv4 addresses to ASes based on the *current* (i.e., 2011-09-01) global routing table snapshot. Then, based on our AS-to-Org mapping results, we group ASes belonging to the same organization and their addresses and thus obtain the address set A for each organization.

Second, for each organization with address set A , we count n_p for each historical snapshot. In particular, we obtain a snapshot of the global routing tables from Route Views every month from 2001 to 2011, and calculate n_p for all snapshots.

Third, we use a simple heuristic to quantify and classify the trend of n_p . Specifically, we use simple linear regression and fit observations over a recent period with a linear model $\hat{n}_p(i) = \beta i + \alpha$, for each observation i in the last M months. Small slopes indicate steady n_p implying a *near constant* number of ASes. A moderate or large positive β indicates growth in the number of key ASes used; that is, a *diversifying* organization. A moderate or large negative slope indicates reduction or *consolidation* in multi-AS use. We estimate the strength of our estimation by summing the residuals:

$$\epsilon = \frac{1}{M-2} \sum_{i=1}^M (\hat{n}_p(i) - n_p(i))^2,$$

where $M - 2$ denotes the degrees of freedom [37]. We consider the trend to be inconsistent if $\epsilon > \epsilon_0$ and large changes are bounded by a constant β_0 . We set $\epsilon_0 = 1$ (i.e., the fluctuation is limited in 1 AS) and $\beta_0 = 0.1$ (i.e., constant usage means less than 1 AS growth/reduction every 10 years). Replacing this simple heuristic with a more rigorous trend analysis is part of our future work.

F.2 Case studies of multi-AS usage

To illustrate how these metrics reflect real-world policies, Figure 12 shows the results of our classification heuristic for two organizations: Google and Comcast.

Both organizations show a number of moribund ASes. The height of each graph is scaled to the number of ASes we discover in the WHOIS data, and the black n_{100} line shows the number of ASes that are routed. The difference shows that each organization has about one-third unrouted ASes (as of 2011-09-01: Google has 8 moribund ASes, or 36%; Comcast, 13, or 27%).

We next describe our classification scheme applied to Google and Comcast. We focus on the past $M = 24$ months (from 2009 to 2011). The trend in multi-AS usage is visualized by the slope of the upper (lower) short light line \hat{n}_{100} (\hat{n}_{80}). As the graph shows, Google exhibits an declining slope (negative *beta*) when using all routed ASes (n_{100}), but a flat slope (small $|\beta|$) when using only the core ASes (n_{80}). This suggests that the core ASes are fairly stable, while smaller and less impor-




Figure 12: The number of ASes (n_p , where $p \in \{100, 80\}$) of Google/Comcast that announce $p\%$ fraction of its addresses, with linear regression \hat{n}_p computed over 24 months. The AS scale extends to how many ASes Google/Comcast has as of 2011-09-01.

tant ASes are being phased out over time. In the case of Comcast, we observe a diversifying trend with respect to all ASes and a consolidating trend with respect to its core ASes for the past two years. Although classified as consolidating, the changes for Comcast’s core ASes are very small (from 6 to 4), suggesting a stable core. The difference between our classification result and the real-world AS usage by Comcast indicates that our classification method may be sensitive to threshold selection.

To evaluate the stability of these trends (consistent and inconsistent), ϵ measures the error against these regressions and is visualized by the deviation of the top dark line n_{100} from the short light line \hat{n}_{100} (or the deviation of the bottom filled curve n_{80} from the short light line \hat{n}_{80}). As Figure 12 shows, neither Google nor Comcast has a significant deviation from the regression line, and thus both are considered as being consistent for the past $M = 24$ months. However, the industry is changing rapidly: if we extend our study back to 2005 ($M = 72$), then none of them is consistent on n_{100} .

Policies in other organizations: To broaden the above examples of policies that affect multi-AS usage, we next briefly summarize our inferences about AS policies for six other organizations.

We see *stable, policy-based ASes in the core* (n_{80} ASes) of many organizations. For example, Figure 2 shows historical routability of all ASes that are part of Google

as of 2011-09-01. ASes are stacked by AS number, with horizontal bars indicating the time periods when ASes are routed, and with darker bars indicating membership in n_{80} . Two ASes have been announcing 80% of Google’s addresses for one year, Google’s main AS (AS15169, AS index: 1) and AS36492 (AS index: 2), designated for WiFi, suggesting a stable routing policy. We see similar patterns for three other large organizations: Figure 13 shows the cores of Verizon (with two geographic one wireless and one access AS), Figure 14 shows Time-Warner Cable (with 6 geographic ASes), and Figure 15 shows China Mobile (with 7 geographic ASes, plus IPv6 and backup core).

Frequently, transient multi-AS usage is the result of *acquisitions followed by AS consolidation*. Continuing with the Google example, in late 2006, Google acquired Youtube (AS36561, AS Index: 16 in Figure 2); the number of addresses announced by this AS have been decreasing since then. In fact, this AS was slowly demoted and then disappeared completely from BGP in April 2011. This suggests that, over time, Google consolidated this service into their core infrastructure. We see similar results for Verizon (see Figure 13) which consolidated ASes from MCI (AS703, AS705, AS3378), an ISP acquired in 2005.

Consolidation also happens when there is a business strategy change. For example, we see *geographic consolidation* in Time Warner Cable caused by an agreement with Comcast in late 2006 [11]. This agreement exchanged subscribers between Time Warner Cable and Comcast to consolidate key regions; it is the likely cause for the death of AS11707, AS13343, AS10311, AS10994, and AS8052 (see Figure 14). These ASes covered areas in Florida, Tennessee and Oregon where Time Warner Cable does not have presence now [5].

Lastly, we found one case where routing policy decisions promote AS diversification: ISC. Although only one AS announces most of ISC’s addresses (AS1280 in Figure 16), we see ISC is using more and more ASes since 2003. Examining these new ASes, we see that each announces a single /24 address block. This policy is consistent with the choice to associate a unique AS with each physical anycast location [18] and with ISC’s operation of the anycasted F-root DNS server. This example illustrates how policy can imply usage of an increasing number of ASes per organization over time, suggesting that this type of multi-AS usage is likely to stay.

We also examined the remaining organizations in our 10 organization list and found very similar behaviors to those shown in Figures 2, 13, 14, 15 and 16.

F.3 Ruling out Churn

While we use n_p to classify organizational use of ASes, this metric focuses only on active ASes. Such a focus

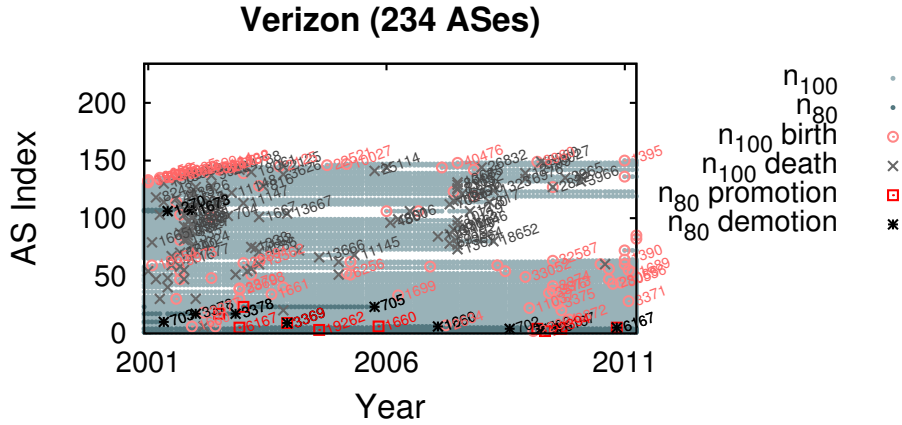


Figure 13: Historical routability of ASes of Verizon.




Figure 14: Historical routability of ASes of Time Warner Cable.

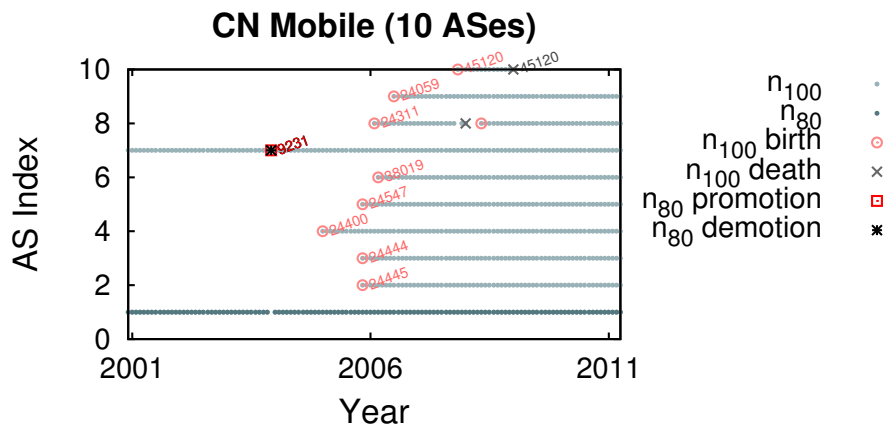


Figure 15: Historical routability of ASes of China Mobile.

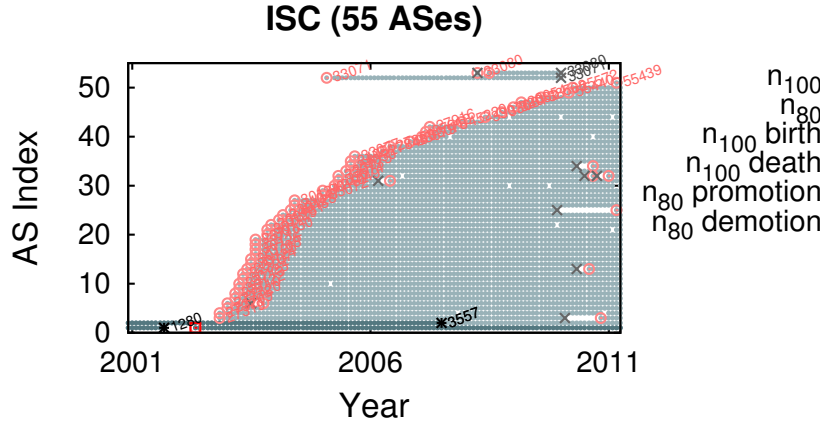


Figure 16: Historical routability of ASes of ISC.

could be misleading for organizations that have *both* growth and reduction in AS usage between two consecutive observations. The churn caused by the simultaneous addition of N ASes and removal of N ASes is not captured by the n_p metric.

To bound the error introduced by AS churn, we measure how often there is an offset in n_p . Since n_p is measured every month, the corresponding offset is given by difference between the number of newly observed ASes and the number of removed ASes in each month. Our examination shows that offsets are rare and small and thus can be ignored. Among 4,388 multi-AS organizations, 3,948 (90%) do not have any offsets at all for all observation intervals. Among the remaining 10% organizations, 423 (9.6%) organizations only have offsets for less than 5% observation intervals, and all offsets are limited in size by 1.

F.4 How persistent is multi-AS usage?

Based on the understanding we gained from these case studies, we next look at all organizations identified by our AS-to-org mapping and present some overall statistics and classification results. Our goal is to answer the following question: Are organizations consolidating or diversifying their use of ASes over time? In other words: Are multi-AS organizations here to stay or are they going away?

We use regression over different durations to see if organizations are consolidating their use of ASes or not. To do this analysis, we begin by selecting all multi-AS organizations. Then for each such multi-AS organization, we perform linear regression on either all of their routed ASes (n_{100}) or on the top ASes that announce 80% of their addresses (n_{80}). We then classify our fit to show AS consolidation, constant use, diversification, or inconsistent trends.

Figure 17 shows these classifications based on time periods from 2001 to 2010 and ending with 2011-09-

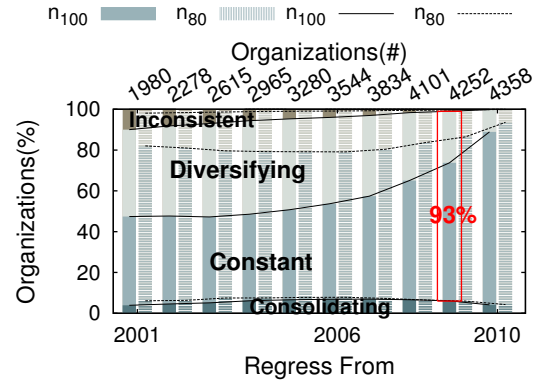


Figure 17: Classification results of multi-AS usage over all multi-AS organizations, based on regression of n_p starting from different years.

01. Solid bars show classification results based on n_{100} while dashed bars are for n_{80} . Solid lines (for n_{100}) and dashed lines (for n_{80}) are used to trace the classification boundaries for clarity. At the top of the figure, we show the absolute total number of organizations that exist at the start of the different regression periods, while the left y axis shows the relative percentage of organizations classified into each category.

Our first observation is that multi-AS use is *not* going away: with a two-year regression based on n_{100} , around 93% organizations are using the same or more ASes. This trend is the same and extends over longer periods if we use n_{80} .

Second, we observe that relatively few organizations are consolidating their AS use over all durations we analyze. While the case studies of our ten organizations show definitive signs of consolidation (3 of 10 are consolidating, and 3 more are inconsistent), when looking across all organizations, we see at most 6% of them are consolidating (for two-year regression based on n_{100}),

with even fewer over all other periods. This difference shows that our selection of the ten organizations is not representative of the Internet as a whole – we chose those ten organizations because of their prominence, large size, and use of many ASes. The vast majority of the multi-AS organizations in the Internet are much smaller, with single-digit numbers of ASes. Big companies typically engage in more acquisitions and mergers and tend to expend more efforts consolidating post-merger.

A third finding is that organizations are much more consistent and constant if we only focus on the top ASes (n_{80}). In fact, we hardly see any inconsistent organizations (top dashed bars), and the percentages of constant organizations far exceed the ones based on n_{100} (compare the dashed line above label “Diversifying” and the solid line above label “Constant”). However, the percentage of consolidating organizations is roughly constant, irrespective of whether we consider all routed ASes or only the top or core ASes. This finding suggests that most consolidations happen in the “core” ASes, while most diversifications occur for the non-core ASes. Presumably, organizations prefer to keep their core small and simple for the ease of management, while they rely on non-core ASes to implement miscellaneous policies.

As the length of regression period (the number of years to look back) decreases, the total number of organizations is increasing (from 1,980 in year 2001 to 4,358 in year 2010), with about 264 new organizations appearing each year. This result reflects the overall growth of the Internet in terms of number and types of Internet-related companies. However, the relative rate of diversification and consolidation appears to have not changed much during these last 10 years. Finally, as expected, with a longer regression period, the number of inconsistent organizations increases – in a dynamic industry, policy changes are frequent.

We thus conclude that the prevalence of multi-AS usage by organizations is persistent and likely to continue in the future.