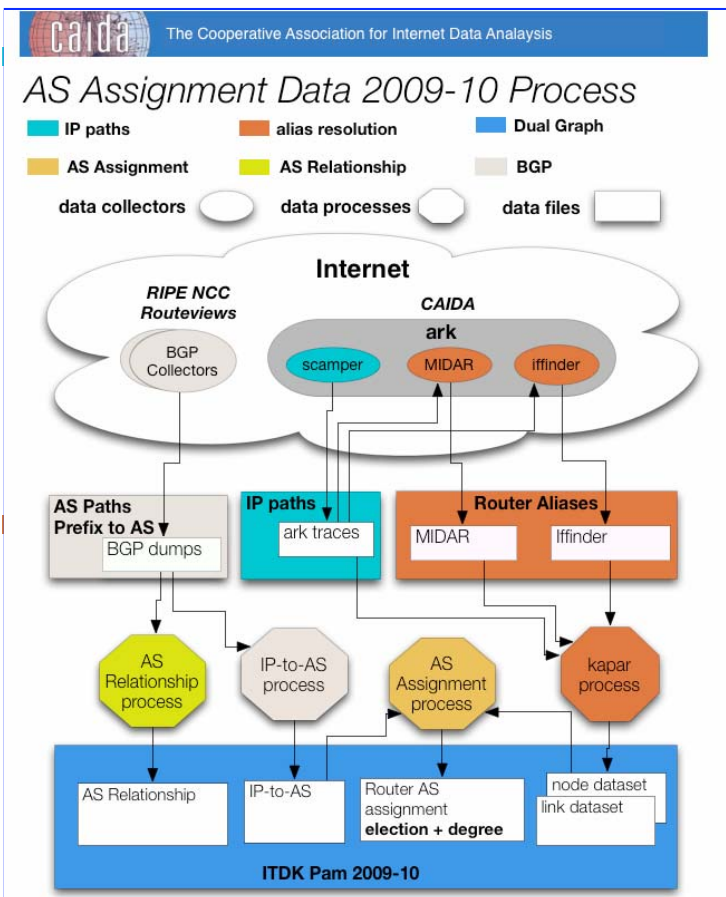




Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers

This document provides supplemental material for the [PAM 2010 paper](#). There is growing scientific interest in the structure and dynamics of Internet topology, primarily at the router and Autonomous System (AS) levels. Substantial progress over the last decade toward understanding and improving the integrity and completeness of router and AS-level topologies separately inspired us to seek a graph construction that merges router and AS-level views of the Internet, i.e., a graph in which routers are annotated with the ASes to which they belong, and inter-AS links are annotated with business relationships. Such a view would capture administrative boundaries and relationships between ASes, while providing sufficient detail about the geography and internal structure of each AS. Inherent limitations and inaccuracies of existing techniques for alias resolution, IP-to-AS mapping, and router-to-AS assignment render this goal challenging. In this work we take initial steps toward merging router and AS-level views into a dual graph representation of the Internet. We start from active measurement (traceroute-like) datasets collected using CAIDA's Archipelago distributed measurement infrastructure (Ark). We then apply state-of-the-art alias resolution techniques to infer which interfaces belong to the same router, creating a router-level Internet map. Finally, we propose heuristics to assign routers to ASes, using information derived from the interfaces that we infer belong to a particular router. This document describes each of the steps in this process and provides links to the resulting dual-graph dataset.



identified 66% of the true aliases from our ground truth dataset, with a 5% false positive rate.

BGP data

To assign IP addresses to ASes, we used publicly available BGP dumps provided by [Routeviews](#) and one of [RIPE NCC's collectors](#) (RCC12) in October 2009. BGP (Border Gateway Protocol) is the protocol for exchanging interdomain routing information among ASes in the Internet. A single origin AS typically announces ("originates") each routable prefix via BGP.

Active measurement data

We collected our active measurements using CAIDA's [Archipelago](#) (Ark) measurement infrastructure. In October 2009 when we gathered the measurements for the PAM2010 paper, the Ark infrastructure had 37 monitors in 28 countries. These monitors used Paris traceroute to probe randomly chosen destinations from each routed /24 prefix seen in BGP dumps from Routeviews over a 28-day collection period in September and October 2009. We call the resulting set of 268 million traceroute paths our *traceroute dataset*, which we use to infer which IP interfaces belong to the same router, a process known as *alias resolution*.

Alias resolution

For alias resolution, we rely on several CAIDA tools: [iffinder](#), [kapor](#), and [MIDAR](#), (recent [tech report](#)). MIDAR (Monotonic ID-based Alias Resolution, a tool we hope to release by April 2010) expands on the IP velocity techniques of [RadarGun](#), while [kapor](#) expands the analytical techniques of [APAR](#). We use the traceroute dataset as input to MIDAR and iffinder, which generate output files used as input to [kapor](#). [kapor](#) heuristically infers the set of interfaces that belong to the same router, and the set of two or more routers on the same "IP link" (a construct that represents either a point-to-point link, or LAN or cloud with multiple attached IP addresses). [kapor](#) correctly

in the BGP tables.

AS relationships

We used the BGP data to annotate each interdomain link with one of three simplified business relationships -- customer-provider (the customer pays the provider), settlement-free peer (typically no money is exchanged), and sibling (both ASes belong to the same organization) -- using the classification algorithm by [Dimitropoulos, et al.](#), resulting in what we call the [AS relationship dataset](#).

AS assignment

The goal of the AS assignment process is to determine the AS that owns each router. For each router r , we first create an *AS frequency matrix* that counts the number of interfaces (known and inferred) from each AS that appears on r . The ASes in this frequency matrix represent the set of possible owner ASes of r . We use the following AS assignment heuristic to assign a router r to an AS.


The **Election** heuristic assigns router r to the AS with the highest frequency in r 's AS frequency matrix. The intuition behind this heuristic is that routers will tend to have more interfaces in the address space of their owner. If two ASes from r 's AS frequency matrix have the same count, then **Election** cannot decide an owner.

The **Customer** heuristic uses the AS relationship dataset to assign relationships to each pair of ASes from r 's AS frequency matrix. **Customer** assigns r to the AS inferred to be a customer of every other AS in r 's AS frequency matrix. This heuristic is based on the common practice that customer and provider routers typically interconnect using addresses from the provider's address space. Consequently, a router with interfaces from both the customer and provider address spaces is assigned to the customer.

For the **Degree** heuristic, we first generate an AS-level graph by assuming full-mesh connectivity among ASes from each router's AS frequency matrix. We then use this graph to generate an AS degree for each AS. **Degree** assigns router r to the smallest-degree AS from r 's AS frequency matrix, i.e., the AS most likely to be the customer AS, based on similar intuition as the **Customer** heuristic.

For the **Neighbor** heuristic, we first determine the set of single-AS routers to which r is connected (its single-AS neighbors). We create a new AS frequency matrix that counts the number of single-AS neighbors of r from each AS. The **Neighbor** heuristic assigns r to the AS with the largest frequency (most single-AS neighbors), based on the intuition that a router is connected to a larger number of single-AS routers in its owner AS. **Neighbor** produces an ambiguous assignment when multiple ASes have the same (highest) frequency.

In case one of the previously described primary heuristics is unable to produce an AS assignment, we attempt to break the tie using one of the other heuristics as a tiebreaker. Our evaluation in the paper shows that **Neighbor** was the best stand-alone heuristic, while **Election+Degree** was the best combination.



PAM 2009-10 dataset

The following datasets constitute the dual-graph representation of the Internet using data collected in October 2009:

- **nodes.20091006.gz (node dataset)** is the set of routers inferred by kapar. For each router, the node dataset contains the set of interfaces that were inferred as being on that router. File format:

```
node <node_id>: <i1> <i2> ... <in>
```

indicates that kapar inferred the node with id `node_id` to have the n interfaces i_1 to i_n

- **links.20091006.gz (link dataset)** contains, for each link, the set of routers and router interfaces that kapar inferred as sharing that link.

File format:

```
link <link_id>: <N1>:i1 <N2>:i2 [<N3>:[i3]] ... [<Nm>:[im]]
```

indicates that kapar found a link `link_id` connecting nodes N_1 to N_m . Note that kapar always determines at least two nodes and the associated interfaces for each link. There may, however, be additional nodes on the same link due to the presence of LANs or clouds with multiple attached interfaces.

By joining the node and link datasets, one can obtain for each router a set of *known interfaces* and *inferred interfaces*. Known interfaces were measured directly and are in the **node dataset**; Inferred interfaces result from entries in the **link dataset** that indicate that router r_1 has a link to interface i_2 on router r_2 , but we did not see an actual interface on router r_1 . The interfaces on an IP link are typically assigned IP addresses from the same prefix, so we assume that router r_1 must have an *inferred interface* from the same prefix as i_2 .

- **prefix2as.20091006.gz (IP-to-AS)** contains the list of IP prefix to origin AS found in the BGP dumps.

File format:

```
<prefix> <length> <origin AS>
```

When multiple ASes were found as the origin of the same prefix, they are concatenated with an underscore.

- **nodes.as.20091006.gz (AS assignment dataset)** contains, for each router found in the node dataset, the AS that is the inferred owner of that router according to our final **Election+Degree** assignment heuristic.

File format:

```
<node_id> <AS> <method>
```

- **as-assignment.tgz (AS assignment code)** contains the code to generate the **AS assignment** dataset.
- **as-rel.a0.01000.20091020.gz (AS relationship dataset)** contains the business relationship associated with each link in the BGP data, computed (see above) using the AS relationships algorithm by [Dimitropoulos, et al.](#)

File format:

```
<AS1> <AS2> <rel>
```

where

rel is -1 if AS₁ is a customer of AS₂

rel is 1 if AS₁ is a provider of AS₂

rel is 0 if AS₁ is a peer of AS₂

rel is 2 if AS₁ is a sibling of AS₂