

The Internet Topology Zoo

Simon Knight, Hung X. Nguyen, Nickolas Falkner, Rhys Bowden, Matthew Roughan
 University of Adelaide, Australia
 {simon.knight,hung.nguyen,nickolas.falkner,rhys.bowden,matthew.roughan}@adelaide.edu.au
 www.topology-zoo.org

Abstract—The study of network topology has attracted a great deal of attention in the last decade, but has been hampered by a lack of accurate data. Existing methods for measuring topology have flaws, and arguments about the importance of these have overshadowed the more interesting questions about network structure. The Internet Topology Zoo is a store of network data created from the information that network operators make public. As such it is the most accurate large-scale collection of network topologies available, and includes meta-data that couldn't have been measured. With this data we can answer questions about network structure with more certainty than ever before — we illustrate its power through a preliminary analysis of the PoP-level topology of over 140 networks. We find a wide range of network designs not conforming as a whole to any obvious model.

I. INTRODUCTION

What is a zoo? The term *zoo* is a common abbreviation for “zoological garden” – a place where animals are kept and exhibited to the public – the earliest of which was that of the London Zoological Society, established 1828 in Regent’s Park. There are many much older collections of animals, but they were described by terms such as menagerie and, apart from the name change, modern zoos differ from a typical menagerie in their guiding principles. The modern zoo is not just concerned with entertainment, but also with conservation, education and scientific research.

The collection described in this paper is a set of network topologies, not animals, so perhaps the term *zoo* (from the Greek *zōon* or “animal”) is inappropriate. However, our goals are those of a modern zoo.

Our Zoo presently consists of over two hundred network topologies from network providers. It is distinguished from other such datasets by the method of collection — we do *not* use an automated procedure such as a traceroute survey as in previous studies (for example, see [1]–[5]). The problems with such surveys are well documented [6], [7]. Here we base our collection on promotional data: maps and other information self published by the owner or manager of the network in question. The results add to the evidence that traceroute is a difficult tool to use for determining topology. For instance, we found one case [5] where the authors declared that a network (Cogent) had 35 PoPs when the network operator themselves advertises 184 PoPs¹.

Natural questions arise concerning the accuracy of our data as well, and we discuss this issue later, but our central argument is that although a published network map may not reflect the exact nature of the underlying network at the current time, it certainly does show the network that the company

intended. Such a map reveals something that no measurement can see: what was in the mind of the network engineer when the network was designed, rather than what was built to meet the realities of day-to-day operation.

The ability to see what a network engineer, manager, or operator believed was important about their network provides insights that traceroute studies lack. Of course, for some purposes a precise map of exactly what currently exists is more useful, and so we see collections of data such as this as complementary to good measurements. In fact, we hope that the Zoo can actually help improve the standard of measured topologies by providing a dataset against which to compare results. However, the network maps we use often provide meta-data about a network that is otherwise unavailable, or at best, subject to large inference errors. For instance we can often see link capacities, node locations, node roles, interconnect locations and so on.

It may be surprising to some that we can collect so many network topologies in this way, but the 200+ networks that we have at present do not even cover all of those that we know (conversion is a time consuming process which is ongoing), and there are no doubt many others that we have not yet found. It seems to be quite common for network operators to publish some form of information about their network. Some even describe their network in legally binding documents such as company reports (e.g., [8]–[13]). Moreover, as networks evolve we expect that companies will update network maps, so we will see the development of these networks. For these reasons, we see the Zoo as an ongoing project. The web page allows for contributions, so in addition to our ongoing efforts we hope to recruit others to add networks to the Zoo.

As with other zoos, there are three primary goals for this collection, starting with scientific research. There is an extensive debate under way on the nature of network topology. On the one hand lie the random graph models (starting with Erdos-Renyi and Gilbert [14] and going forward through Waxman [15], and more recently power-law graphs [16]–[19]). On the other hand lie “designed networks” such as the structured networks of GT-ITM [20], [21] or HOT (Highly Optimized Tolerance) graphs [22]. Proponents of power-law and HOT graphs seem convincing, but both are hampered by lack of accurate data. In the few cases where a commercial network has been used the data have not been published. This lack of accurate, public data has been a severe constraint, preventing performance of repeatable scientific research. One may reasonably argue that scientific progress cannot be made in this area without an accurate, public set of data.

In addition to scientific interest in the very nature of these networks, the provision of network datasets in an easily

¹http://www.cogentco.com/us/network_map.php

accessed format may prove useful to the broader network community: network topologies are necessary for testing many networking algorithms. At present, very few are available apart from those based on traceroute studies.

The Zoo’s second goal is education, though this is obviously related to the research value of the data. We will learn lessons from this research, and those lessons will help educate researchers². The Zoo will also provide a set of network examples to help educate the next generation of network engineers. Most current examples are contrived. Real examples are more compelling.

The third goal of the Zoo is conservation, though not in the ecological sense. The maps and other promotional materials that we use here are ephemeral. Once their use-by date is exceeded they are removed from public view, but more than that many companies do not make any attempt to store historical archives of their network designs. Apart from scientific uses, such data may even be useful to those companies at some point in the future to understand the development of their own network. The Zoo currently includes a number of historical network maps, going back to those of the ARPANET, published by Cerf and Kahn [23]. The Zoo contains several other examples for which we have multiple views of that network’s development over time.

This paper is more than a description and classification of the Zoo itself. We also present, in Section IV, some preliminary analysis of these networks. We focus in this paper on maps at the Point of Presence (PoP) level, which are interesting because this level relates to the network design problem. It is also the level which concerns peers and customers as it determines where they can connect to the network, and it’s also the level at which reliability and redundancy are often considered. What we see in the 141 networks studied is that there is no “one true network model”. There are a very wide range of networks ranging from hub-and-spoke networks, to trees, to more highly connected graphs. However, we do observe some trends. In particular we see more hierarchy with increasing network size. We also make one new observation that the neighborhood of a PoP seems to be limited to about 20 or so other PoPs. There is no physical or technical constraint that enforces this at the PoP level, so it will be interesting in the future to explore the reason for the presence of this limit.

II. THE DATA

Before we begin discussing the details of our topological data, let us first define our terminology precisely as topology is a woefully abused term. We define *topology* to mean an undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, which abstracts the connectivity of a data communications network. In fact, we really mean a multigraph, as multiple edges are allowed between a single pair of nodes (formally, \mathcal{E} is a multiset).

Care must be taken to define the nature of the nodes and edges of the graph. Internet topologies have been given for each of the seven OSI layers, for instance edges may refer to physical cables, virtual network layer connections, or

even the HTML links between WWW pages. Other types of topology are also possible, such as those reflecting hierarchical approximations, say by combining some groups of routers into Autonomous Systems (ASs) or Points-of-Presence (PoPs). The Zoo contains topologies of various levels of detail, from physical fiber, through to virtual/logical connectivity between ASs. We admit most types of networks to the Zoo, but ensure that in the data the type of nodes and links are precisely specified.

A. Measurements

There are various strategies available for measuring network topology. The most direct way is to ask the network itself. IP routers are managed through *configuration files* describing the current operation of the router, and which can be used to measure a network [24]³. Precisely because of the quality of information contained in these files they are considered highly sensitive and are rarely allowed outside an organization. Such data may be used to construct the type of map we use here, but is otherwise rarely available to researchers.

The second class of techniques involve IP-level hacks that ideally return the path between two points. The IP header option field “record route” [26], [27] returns the route of a packet as it traverses the network, but is often not enabled due to security and performance concerns. The more common approach is *traceroute* [28], [29]. Despite being commonly used, *traceroute* has many well-known deficiencies (summarized in [6], [7]). There are nevertheless many studies of network topology using traceroutes (for examples see [1]–[5], [30]), but the resulting network topologies are not very accurate [6], [7]. Moreover, verification of these topologies is made difficult by lack of ground-truth. One of the potential uses of the Zoo data is to establish ground-truth data to use in improving measurement-based approaches, which can in principle survey a wide range of networks.

We performed comparisons between our dataset and one of the most recent and advanced traceroute based methods and found large differences. For instance, the example of Cogent’s network in [5]. There are several possible causes for this discrepancy, the most likely being traceroute measurement errors, differing definitions of a PoP, and differing network boundaries. However, we maintain that a network operator is in a better position to define details such as the edge of its network, and so their view should take primacy.

The third group of strategies for topology inference is based on network tomography. The statistical nature of these approaches again leads to errors.

Instead of the existing automated methods we adopt here a simple, manual approach. Many companies present public material about their network, primarily for promotional purposes. They wish to sell their network. We capture this information, and manually transcribe it into a common data format – in the following section we describe in detail our process of capture and analysis for these datasets.

²In fact, some of what we learn from this data is already accepted by network engineers, and so part of the value of this data is in educating researchers. However, it is dangerous to be too trusting of received wisdom. We should maintain skepticism, and verify even that which is “well-known”.

³A related approach is to use a routing monitor (e.g., [25]), which observes routing protocols and uses this information to construct a network topology, but this also requires privileged access to the network in question.

B. Data Collection and Formatting

Some network operators provide a piece by piece description of their network, but the most common form of published information available to us is a network map. Such maps often show PoPs and their interconnects, but may provide much more detail. Some care goes into such maps because they are a form of advertisement and therefore have legal requirements for accuracy; they are highly visible to potential customers; and finally, network engineers are often proud of their work, and many would very much like to display it at its best.

Often these maps are simple images, but in some cases they are interactive maps (for example using Flash). In the case of images, we manually download the map and then transcribe it into an annotated graph. Dynamic maps are more difficult, and often require several passes to zoom in on details and transcribe. Supplemental data in the form of equipment registers or other descriptions of the network are used where available to label links and nodes. We have collected over 200 such maps and associated data, and make no claim that we have an exhaustive list.

Maps are converted using yEd [31], a freeware tool for use in graph-drawing. It allows us to trace the network elements such as routers and PoPs directly overlaid on the map itself, with annotations such as node names and edge capacities. A graphical diagram editor speeds up the tracing and annotation steps, and reduces errors by allowing a visual comparison of the original source image and the transcribed network.

Once we are satisfied with the transcribed network, we export the topology into GML (Graph Markup Language) format. yEd supports a number of graph formats (for instance GML and GraphML), but GML meets our immediate needs for a flexible, easily readable format. We wanted a format that was easily computer readable, but also human readable. A graph can be most easily understood pictorially, but it enhances our ability to double-check data if we can read it without the intervention of third-party software.

We also needed a graph format that was extensible. Different network operators provide different information in their network descriptions. Some may provide PoP-level or router-level maps, or detailed information about the physical media used for a link, while others may show links that are planned for the future. We don't know all of the data that we might need to store, and so we need a format that can be extended.

Binary file formats are compact, but are difficult to read and to extend. Adjacency matrices capture the graph's structure but have limited scope for storing attributes such as node names.

GML [32] is a simple, text-based format. It simply lists nodes and edges, with extensions to allow node and edge attributes to be stored. Attribute information is represented inside square brackets as key-value pairs. GML is also supported by a number of tools, and easily ported to other formats (we provide simple scripts to do so). Hence we use GML as our core file format. We provide a simple network example in Figure 1, to illustrate the data format.

However, we understand that other users of this data may find other formats more convenient. XML-based languages such as GraphML [33] are easily parsed by machines — XML processing libraries exist for most popular programming languages, making it simple to work with data from the

Zoo. XML is also extensible by design, allowing it to handle arbitrary attributes. We use GML as our core file format, but provide the data in GraphML format as well. We also provide scripts to read and convert the graph data into other formats such as a simplified adjacency matrix representation.

One of the major advantages of GML is that it can be read by NetworkX [34], an open-source graph library for the Python programming language. NetworkX is fast, well supported and includes many graph analysis algorithms. It is these we use to perform much of the statistical analysis presented later.

C. Meta-Data

One of the chief advantages of our approach is that many network maps contain additional data. We include such meta-data in the records, for instance:

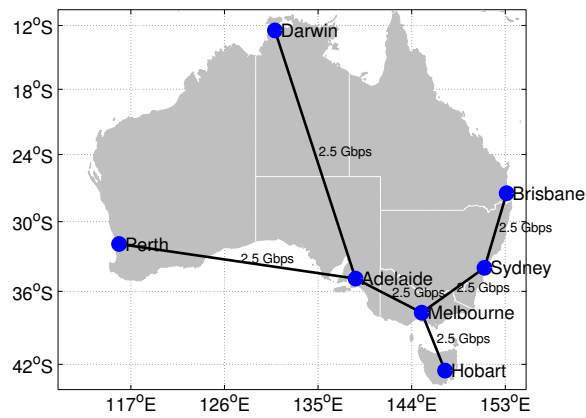
- link types or speeds;
- longitudes and latitudes of nodes obtained through geocoding of PoP locations;
- a URL (Uniform Resource Locator) showing where the data was obtained;
- the date-of-record, i.e., the date that the map was representative of the network (in cases where the network map was dynamically generated we record that);
- we also record the date we obtained the network map;
- a classification of the type of network. This last point requires much more discussion and we will do so in Section III;
- a link to other related networks.

D. Accuracy

How accurate is the Zoo's data? The maps are created by network companies themselves, so they are directly based on ground truth. However, some network operators clearly produce these maps manually, potentially leading to inaccuracies in their depiction of their own network. There are two reasons that these errors are less significant than those of prior studies.

- The network maps we use are all public documents, and so must satisfy standard due diligence requirements for an advertisement or official corporate publication. That is not to say that all corporations are perfect — it is easy to make mistakes in drawing the map — but a network operator is unlikely to publish a worse map than the one they use in their own network operations.
- Some network maps may idealize the network. However, we argue that in these cases, we are seeing what was in the mind of the network engineer when the network was designed. In this sense, the idealized view of the network may be more interesting than its implementation (though for some purposes it may be preferable to see exactly what was implemented).

On the other hand, network operators do perform simplifications in some cases, most notable, many of the maps report PoP-level, not router-level topologies. The datasets include the level at which they are defined, and it is important to be aware of this issue when using this data for research. For a very simple instance, consider a network reliability study. A single PoP may consist of a number of redundant routers, so the



(a) Network aus_simple.

(b) GML

```
graph [
  directed 1
  GeoLocation "Australia"
  GeoExtent "Country"
  Network "aus_simple"
  Classification "Backbone, Transit"
  Creator "Topology Zoo Toolset"
  LastAccess "03/12/10"
  Layer "IP"
  label "aus_simple"
  Source "Example"
  Version "1.0"
  Developed "Developed"
  NetworkDate ""
  hierarchic 1
  Type "COM"
  DateObtained ""
  node [
    id 0
    label "Perth"
    Internal 1
  ]
  node [
    id 1
    label "Adelaide"
    Internal 1
  ]
  .
  .
  edge [
    source 0
    target 1
    Speed "2.5"
    Label "2.5 Gbps, Ethernet"
    Units "G"
    Type "Ethernet"
  ]
  .
  .
]
```

(c) GraphML

```
<?xml version="1.0" encoding="utf-8"?>
<graphml xmlns="http://graphml.graphdrawing.org/xmlns"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://graphml.graphdrawing.org/xmlns/1.0/graphml.
    xsd">
  <key attr.name="key" attr.type="int" for="edge" id="d23" />
  <key attr.name="Label" attr.type="string" for="edge" id="d22" />
  <key attr.name="Speed" attr.type="string" for="edge" id="d21" />
  .
  .
  <graph edgedefault="undirected" id="">
    <data key="d0" />
    <data key="d1">Australia</data>
    <data key="d2">Country</data>
    <data key="d3">aus_simple</data>
    .
    .
    <data key="d7">COM</data>
    <data key="d8">IP</data>
    <data key="d9" />
    <data key="d10">Example</data>
    <data key="d11">1.0</data>
    <data key="d12">aus_simple</data>
    <data key="d13">1</data>
    <data key="d14">1</data>
    <data key="d15">03/12/10</data>
    <node id="0">
      <data key="d16">1</data>
      <data key="d17">0</data>
      <data key="d18">Perth</data>
    </node>
    <node id="1">
      <data key="d16">1</data>
      <data key="d17">1</data>
      <data key="d18">Adelaide</data>
    </node>
    .
    .
    <edge source="0" target="1">
      <data key="d19">G</data>
      <data key="d20">Ethernet</data>
      <data key="d21">2.5</data>
      <data key="d22">2.5 Gbps, Ethernet</data>
      <data key="d23">0</data>
    </edge>
    .
    .
  </graph>
</graphml>
```

Fig. 1: Example of GML and GraphML file formats. Many of the tags will be explained in the following section.

likelihood of the whole PoP failing is much smaller than for a node in a router-level graph.

Most of the maps in the Zoo come directly from the network operators, but some have been derived from secondary sources. We don't wish to exclude any interesting data from the Zoo, however, in these cases, the data is potentially less reliable. Hence, we include in the data a "provenance" field taking the form: *primary* (meaning it comes from the operator itself), *secondary* (from a reputable secondary source, for instance the scientific literature) or *unknown*. For studies requiring accurate maps we suggest that only data with primary provenance are used.

A second question of accuracy is "How accurate are our transcriptions?" We have transcribed a large number of maps so it is inevitable that some errors occur. However, we have tried to minimize errors by (i) using a graphical tool so that the transcription process is closely matched to the maps, and (ii) making sure that each network is transcribed by one person, and then checked by at least one other person.

Despite this care, there are still difficulties in interpreting some maps. The most pernicious problem is links that join without a node. There are two possible explanations for such joins: (i) that there is really a three way link, and that the correct graph representation is to join three nodes (with three links), or (ii) that there is a *y*-junction, and one node is connected to the two others, but there is no third link. However, we do not know which is the reality, and so we introduce a "blank" node at the join. This is the biggest source of potential inaccuracy in the Zoo, but these nodes are flagged in the data, and so it is possible to take the appropriate care to eliminate problems caused by such ambiguities.

Another potential source of error is a network where it's too complicated to follow the tangle of links, or where it's unclear whether some nodes are real or logical. In such cases, we exclude that network from the Zoo.

Ongoing quality control is an important part of this project, and the web page also has links to a discussion forum, to allow ongoing contributions to the accuracy of the dataset. The forum provides a way for users of the data to point out flaws, either in transcription or interpretation of the data. The ongoing improvement of the quality of the Zoo's data is as important as care in the initial collection.

E. Visiting the Zoo

The data is stored at www.topology-zoo.org. It is viewable through a table containing meta-data about the networks, or as a single archive file. Scripts are provided for easy access and translation of the data. An addition goal of the web page is to allow contributions to the Zoo from third parties.

The data at present consists only of data provided by operators, however, we see no reason in principle why we could not include crawled topologies, for instance, social network topologies. Obviously such data would need to be classified and tagged appropriately, and details of the data's limitations published. Further, we require that additional datasets conform to the same data format, though writing translation scripts is not usually difficult (the tools we currently provide include a translation script for Rocketfuel data [1]). GML is flexible enough to allow for such extension.

Finally, because the dataset represents a growing collection, we use a version control system to keep track of the state. Archives of the dataset at particular time points will be kept for comparisons with past studies. We ask of any researchers who make use of this data that, apart from taking care to first understand the limitations of the data as documented above, they cite the exact version of the dataset they use.

III. CLASSIFICATION

Having collected a number of inhabitants for our Zoo, we are left with the question of how to classify each of our new inhabitants. Classification of species or entities is important for several reasons. First, we wish to describe what the Zoo contains. Second, we suspect that different types of networks will have different qualities, and we wish to test that hypothesis. Third, we can now identify when we have discovered something new and hence extend our classification.

There have been a number of prior works on classifying networks, in particular, Autonomous Systems (ASs) [35], [36]. The focus of these has been machine learning techniques applicable to classification of all the tens of thousands of ASs in the Internet. Here our focus on a smaller subset allows us to classify the networks manually. We also have a more detailed data source in the information a network operator provides.

Past classification efforts [35], [36] have tried to cluster networks into disjoint categories. We could easily extend this into tree-like classification resembling Linnaen taxonomy of living organisms [37]. However, the tree-like classification popularized by Linnaeus in his *Systema Naturae* [37] is hard to apply. Even in biology where the tree based taxonomy was later ratified by the theory of evolution there are many cases where purely description taxonomy fails to identify the evolutionary tree (for instance compare anteaters in Australia and South America. Both have similar adaptations for their exclusive diet of ants and termites, but the Australian Anteater or Echidna is not even a placental mammal – it lays eggs), and Linnaen taxonomy completely fails at corner cases such as the Duck-billed Platypus (this and the Echidna are the only Monotremes). We see similar problems in the existing work on network classifications, and here there is, as yet, no underlying tree-structure to justify a tree-like classification system.

We must remember that Linnaeus' system was not originally proposed as a true categorization of natural groups, but was to provide clear identification. The initial point of such classification is to simply list features of organisms, and it is this approach we adopt. We create a set of binary classifications describing whether a network has a particular feature, but these features are not disjoint. The advantage is that we can easily handle corner cases that would be problematic for an exclusive classification scheme, and do so with far fewer classes.

Classification *tags* have the additional advantage that it is easy to add new types without changing the existing classifications, something that would be impractical with a disjoint class model. This has proved useful as we learnt of new classifications that could be usefully added to the data. The classifications we have added so far are described below.

A. Type

At the most basic level we classify our networks as Commercial (COM) or Research and Education Networks (REN). So far we have transcribed 112 COM and 120 REN networks.

Our secondary type classification is related to the role the network plays: *backbone*, *testbed*, *customer*, *transit*, *access* and *internet exchange*. We have only associated these tags with networks where we have clear evidence that the tag should be employed and therefore we believe our false positive rate is low, but concede that our false negative rate will only be established if and when the network owners comment on our classification. However, we choose our categorization based on features that are desirable for a company to advertise (for instance IXPs must self-identify in order to advertise), and so we believe that the false negative rate is acceptable.

The *access* tag denotes a network that provides edge access (for instance via DSL, dialup, or fiber) to individuals.

The *backbone* tag denotes a network that connects at least two cities⁴. Most such networks self-identify as backbones as this is often a key feature to advertise. The smallest backbone considered contained only three nodes.

The *customer* tag is used when a network provided a higher level of services to its customers⁵ than transit. We classified a network with this tag if the services provided required per-customer state: for instance, web hosting or electronic mail. With the introduction of per-customer state, the provider must have a customer service model that is not driven purely by the technical requirements of maintaining connectivity and core services (DNS, routing, etc.). This tag is applied when a provider clearly advertises a web-hosting, e-mail or co-location facility, or similar per-customer state service, for their connected organizations.

The *testbed* tag denotes a network with support and facilities for experimental network protocols or hardware.

The *transit* tag is used when an operator indicates that the network is connected to other networks in a way that supports transit, and the operator indicates that they have the infrastructure required to function at this level. In order to apply this tag, we required clear evidence in the web pages and network diagrams that transit was not only possible but was currently being provided.

The *internet exchange* (IXP) tag is used where the network in question exists only as a nexus of other networks.

The other major aspect of network type is the layer of the network. We provide tags indicating the layer (1-3) and more information about the type of technology/protocol being used where available. Within the *IP* networks we distinguish router-level graphs from PoP-level graphs (though in many cases the PoP-level graph can be derived from the router-level). We concentrate, in Section IV, on PoP-level graphs, of which we have 141.

⁴The backbone tag appears redundant in this paper because in the PoP focus of our later investigation naturally focussing us on networks that all have this tag, but note that the Zoo itself contains a wider variety of networks and so the classification is important.

⁵In many cases customers might not be individual users, they may be businesses or research organizations.

B. Range

As for any species, to understand its nature, we must understand the range over which it operates. An animal with a wide-ranging habitat will have a far greater influence on other species. Similarly, we may more accurately compare networks if we focus on their area of influence [38]. The tags for this categorization are taken from the set *metro*, *region*, *country*, *country+*, *continent*, *continent+* and *global*.

A *metro* network is one that spans a city, or a city-sized area. Likewise for the *country* and *continent* designations. In each of these cases we also add a tag describing the range, e.g., in the case of *continent*, it designates the “continent”: North America, Europe, Asia-Pacific, Latin America, and Africa.

A *region* network is approximately the size of a state or a small number of states, where the number of states involved is not a substantial part of the containing country. The difference between the size of states in Australia, which can be as big as countries in Europe, demonstrates the need for a flexible definition of the “greater-than-metro but smaller-than-country” range classification.

The *country+* (and *continent+*) classifications are used when the network is mostly located within one country (or continent) but has points of presence in another that do not correspond to a significant number of the total PoPs. The label is needed because there are many networks that are easily identified as belonging to a country (or continental) region, but for expedience have one or more routers outside the country.

Where a network has significant presence in at least two continents, it is labelled as a *global* network.

IV. POP-LEVEL STATISTICAL ANALYSIS

We have discussed various approaches for estimating network topology. The advantage of automated measurement-based approaches is that they can survey large numbers of networks, and therefore potentially see a larger proportion of the Internet, and examine it over time. However, the unknown quality of the data, and the known inaccuracies make it unsuitable for some purposes. Our approach is manually intensive, and relies on public data. It is unlikely we will see more than a few thousand networks displayed in this way, or have the resources to parse more than this number. It could be argued that it is preferable to base a study on a small number of accurate network topologies, even with potential sampling biases, than 30,000 measurement artifacts, though in fact the two sets of data are complementary, and each provide their own perspective.

So what do we see in this data? In this section we perform some preliminary statistical analysis of the networks we have collected so far. Here we will focus on the PoP (Point-of-Presence) view of the network. These are interesting as a group, but more importantly it would be misleading to compare apples with oranges, so we want to ensure that we present a comparable group of networks.

A PoP can be roughly defined [5] as “a group of routers which belong to a single AS and are physically located at the same building or campus.” A single PoP often offers Internet access over a much wider area, so it can often be conflated with a metropolitan area or some equivalent. The PoP-level

view is interesting from several points of view: it shows the wide-area links, which are most interesting when it comes to network design optimizations; it concerns peers and customers as it determines where they can connect to the network⁶; and it’s the level at which reliability and redundancy are often considered. The interest in networks at this level is reflected by the fact that it is the most common form of published network map. Router-level map may be seen as too proprietary for publication, but more often it is probably felt that such maps contain many implementation details that aren’t interesting to anyone except the network’s engineers. In this section we will focus on the Zoo’s PoP-level maps.

We do not claim these networks are a random sample of the Internet. There are biases in the collection methodology in that we can only transcribe published networks. So the results here are descriptive rather than representative, but nevertheless provide some insights.

A. The Networks

Of the 232 in the Zoo, 141 are at the PoP level, and we denote these the Networks Under Study (NUS)⁷. They are split between 59 REN and 82 COM networks. There are many more commercial networks in the Internet, but academic networks are often free of the commercial considerations that lead to network operators hiding their topology, and so there is a natural bias towards these networks in our sample. 106 of the NUS are from a primary source, 24 from secondary sources, and only 1 has an unknown provenance.

Tables I and II show the types and ranges of the NUS, and in brackets the whole Zoo. In this study we have no IXPs (these networks are uninteresting from a PoP-level perspective), and all networks are classified as backbones (because they link multiple PoPs). The other classification tags are more useful in subdividing the networks. We suspect these proportions are also biased as compared to the actual proportions in the Internet, but again less so than most datasets. For instance, most measurement techniques are poor at seeing the “edge” of the Internet. That is, they have trouble seeing the topology of the (sometimes) small providers who connect users to the Internet. This is a result of the fact that traceroute surveys can only see forward paths, and so can only see the topology at the edge if there are local traceroute servers, whereas most such servers tend to be located towards the “center” in transit providers. Thus, while the sample we have may not be statistically unbiased, it presents a clearer view of the wide range of the behaviors that are possible.

Table III reports the dates-of-record for the NUS (and in brackets the wider Zoo). Some of the older networks may no longer be operational, or at the very least will have changed, but conservation of a historical view of networks is part of the goal of this project.

The other aspect that is interesting is the size, that is, the number of PoPs in the networks. A reader might suspect that there is a bias in the NUS towards smaller, more easily

⁶The Zoo data often includes details such as the links to external networks, however, the networks under study in this section will only be comprised of PoPs, i.e., external links will not be included.

⁷Dataset Zoo.PoP-level.v1.00.

type	COM		REN		Total	
access	19	(26)	0	(0)	19	(26)
backbone	76	(105)	59	(119)	135	(224)
customer	61	(84)	24	(34)	85	(118)
testbed	4	(4)	9	(19)	13	(23)
transit	54	(69)	12	(41)	66	(110)
IXP	0	(0)	0	(4)	0	(4)

TABLE I: Types of the NUS. The number in brackets is the number of networks present in the complete Zoo.

Range	NUS	Zoo
Metro	1	2
Region	5	23
Country	101	161
Country+	6	7
Continent	16	27
Continent+	1	1
Global	10	10
Unclassified	1	1
Total	141	232

TABLE II: Range of the NUS, and the rest of the Zoo.

<2000	2000-2005	2006-2010	Current	Dynamic
1 (16)	9 (25)	26 (58)	93 (117)	12 (16)

TABLE III: Date-of-record for NUS (Zoo). Current means the map is provided without a timestamp but was listed on the operator’s website at the time it was downloaded. Dynamic maps are generated on-line.

transcribed networks. We have tried to avoid this particular bias. Figure 2 shows the cumulative distribution of the NUS by size. We can see that despite these being PoP-level networks, some are still quite substantial. The Zoo actually contains a few even larger networks (one with 751 nodes) but these are optical fiber networks not included in the statistics below. We have also made some effort to make this a fair comparison. Not all network maps use the same definition of PoP. We found three REN networks which used the term PoP for university campuses connected to their network. In these cases we remove these “edge” nodes from the graph in the reported statistics, but show the affect of this in Table IV (for size and average node degree).

Network	Size		Average Degree	
	Before	After	Before	After
ULAKNET	79	3	2	2
PERN	127	12	2.03	2.33
Sinet	74	13	2.05	2.31

TABLE IV: Networks which have “edge” nodes. Before and after trimming the edge nodes.

The other key aspect of the network concerns the edges. We could measure the total number of edges, but it is more telling to calculate the average node degree ($2 \times$ number of edges/number of nodes) in the network. Figure 3 shows the distribution of the average node degree over all networks. Typical values lie in the range 2-3, which is fairly low. The lower bound for average node degree for a connected network with n nodes is $2 - 2/n$ for a tree. One explanation for the low node degrees would be if the networks were disconnected. In fact there are 17 disconnected NUS. In 8 cases the disconnected segment is a single node, and in all

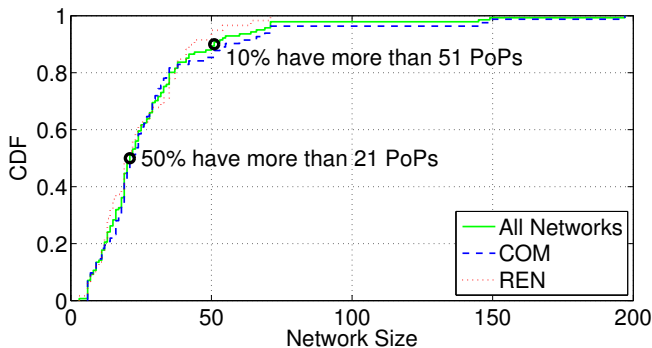


Fig. 2: Cumulative distribution function of the numbers of PoPs (the network size) in the NUS. The solid line shows the distribution for all networks, and the dashed lines show the COM and REN networks separately. They are all very close.

but 2, the disconnected nodes are single nodes. Only two, show more than one non-trivial connected component. In Figure 3 (and all subsequent graphs) we only consider the largest connected component of these graphs.

The question remains then “Why are the average node degrees so low?” These are PoP-level graphs, so there is no physical or technological constraint on node degree. We shall discuss this question further in the more detailed statistical analysis to follow.




Fig. 3: Histogram of average node degree. The average node degree varies from 1.66 to 4.5.

B. Graph Metrics

The preliminary statistics discussed above were primarily to help describe the NUS. In this section we start a more detailed analysis of these networks.

We start by looking for simple scaling relationships between network size and other network variates to determine if the type of the network has any influence on such scaling. Figure 4 shows the average node degree versus network sizes. We also plot the lower bound on average node degree for a connected graph: $2 - 2/n$. In many respects this supports the inference from Figure 3 that node degrees are low. A few cases are on the bound indicating the networks are trees. Apart from the lower bound there is no strong relationship between node degree and network size. Some models of network growth would predict an increase in average node degree as a network grows, but

we don’t see that. We do seem to see higher degree networks amongst the COM networks – future work will be aimed at determining if this is a statistically significant difference.




Fig. 4: Average node degree vs network size.

Another question that has been frequently debated is whether networks have power-law degree distribution [7], [16]. This debate has been hindered by poor definition of the topologies under consideration, and poor data. Even here, we cannot answer this question with any certainty. However, this is not a problem with the Zoo data, but more a problem with the concept – the concept of power-law degree requires large networks so that a large range of degrees can be measured. PoP level networks simply don’t have enough nodes. It’s hard to imagine a network operator with the tens of thousands of PoPs that would be needed.

Nevertheless, we can consider the merits of the idea underlying the power-law degree distribution, that of a highly varying distribution. As noted, we have only a few measurements for any one network, so looking at the distribution won’t help, but we can look at summary statistics such as the coefficient of variation of the node degree. We do so in Figure 5 (in comparison to network size). Again we see little evidence for any systematic relationship between the statistic and network size or type. However, tellingly, we do see that these values are distributed around one, and all are below two. Simple simulations from a standard highly varying distribution (the Pareto distribution with $\alpha = 1.5$) also generate some cases around this value, but we would expect to see some cases with much higher coefficients of variation. The measured coefficients are inconsistent with a highly-varying distribution.




Fig. 5: Coefficient of variation of node degree.

Another useful graph metric is *assortativity* [39], which

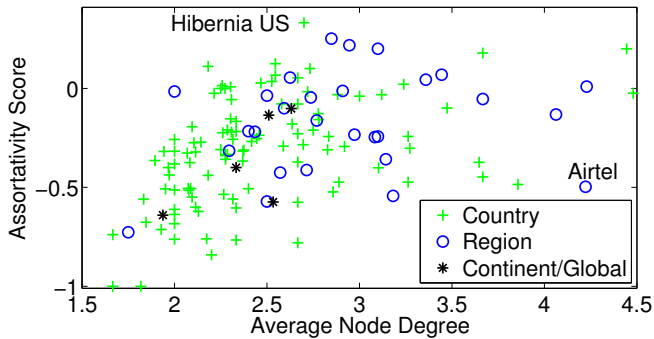


Fig. 6: Assortativity mixing score. Most networks have negative scores, with a few slightly positive.




Fig. 7: Example summary statistics on three simple networks.

measures the mixing properties of the nodes. Assortativity refers to a preference for a network’s nodes to attach to others that are similar or different in some way, here we measure with respect to node degree. Positively assortative networks are those where high degree nodes tend to connect to high degree nodes and low degree nodes connect to low degree nodes. Negatively assortative networks reverse this relationship. We plot the assortativity of our networks in Figure 6. Most assortativity values are below, or only slightly above zero, meaning that high-degree nodes tend to connect to lower-degree nodes, not each other.

Examined in more detail the figure also shows an apparent trend to higher, but still negative, assortativity values for networks with higher node degree, though not clearly correlated with other measures of size such as geographic extent.

The presence of a hub-and-spoke network brings up the question of how common this design or other similar designs are. The assortativity suggests that other pure hub-and-spoke networks are rare, but a superficial look over the data found several others that were close. We can measure the degree to which a network is “hub-and-spoke” like by looking at the correlation between node degree and *closeness centrality*, which is defined [39] for a vertex v as the reciprocal of the sum of geodesic distances to all other vertices of the graph. In a hub-and-spoke network, the “central” node will have high closeness to all other nodes and high degree, and hence a high correlation between degree and closeness centrality. Figure 8 shows the correlation coefficients of the two metrics across each network as a function of network size. We see that quite a few networks have high correlation – corresponding to being somewhat hub-and-spoke like – but that this decreases for larger networks. Moreover, we see some differences between




Fig. 8: Correlation between closeness centrality and node degree. High correlation means high degree nodes are at the “center” of the networks. The correlation decreases as network size increases. Smaller networks seem to have high degree nodes at the core and larger networks show less such structure.




Fig. 9: Size of the neighborhood of the largest node. There appears to be a maximum neighborhood size just above 20.

networks from different regions. Asia-Pacific networks seem to be more hub-and-spoke like than American or European networks, though as with other plots there is considerable overlap between classes.

A second way of examining the same issues of hub-and-spokiness is to look at the size of the neighborhood of the node with the largest degree (the neighborhood is the set of adjacent nodes). Figure 9 shows this metric, and we see again a number of graphs where the neighborhood of the largest node encompasses much of the network. However, this graph also suggests another interesting phenomenon. We see a maximum neighborhood size of about 23 PoPs. There are no physical limits to the neighborhood of a PoP because we may use multiple routers to overcome technical limitations such as port density or throughput limitations. Therefore the reason for this limit must originate elsewhere. We cannot explain it at the moment except to suggest that a “mega-PoP” is difficult to manage. The complexity of this PoP becomes such that it is just easier to break the network into two or more core PoPs, and that also mitigates the potential damage of a PoP failure resulting from say a fire or major power-outage.

This in turn suggests the creation of hierarchy in larger networks. A graph metric that we can use to start to examine hierarchy is *betweenness*, defined for an edge as the number of shortest paths that traverse that edge [39]. The number




Fig. 10: Normalized edge betweenness as a measure of hierarchy.

of paths in a network grows as $n(n-1)$, so to perform fair comparisons we normalize our betweenness values by dividing the maximum by the average for a network, and we plot the results in Figure 10. We see a quite distinct trend of increasing betweenness with network size. We hypothesize that this results from an increasing degree of hierarchy in networks as they increase in size. As a network grows, it often takes on some elements of hierarchy, and these are reflected in some links becoming trunks between regions with a higher betweenness than the average links in the network. On Figure 10 we also plot the betweenness values for a balanced tree. We can see that it loosely supports the same trend we see in the data.

Another feature we noted from superficial examination of the NUS was that some appeared to have a “core” of densely connected hub nodes, with many low-degree nodes hanging from this core. To examine this more formally we consider the size of the *2-core* of the network. The *k-core* of a network is the subgraph that contain only nodes of degree k or above [39], so the 2-core excludes all degree 1 nodes (this must be performed iteratively as some nodes become degree 1 once the first set of degree 1 nodes are removed). Figure 11 shows the ratio of the size of the 2-core of a network to its size. We see the tree-like networks clustered at 0, and a general increase in the size of the 2-core with average node degree. We also see some separation between REN and COM networks. However, the most interesting detail of this graph (as with many others) is the variety – the proportion of the network covered by the 2-core runs from none of the nodes, to all of them with a selection of values in between. Although there is a correlation with average node degree, this provides a lower-bound on the proportion rather than providing a clear trend.

Summary: The picture we get of our PoP-level maps is mixed. We see evidence for or against many phenomena, for instance:

- against power-law degree distributions;
- for hub-and-spoke like behavior;
- for hierarchy;

but the evidence is never completely convincing, reflecting the sheer variety in the networks. If there is any message in this data, it is that there are as many types of networks as there are network designers.

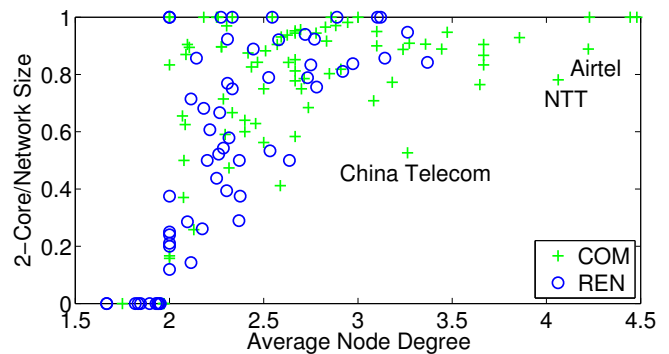


Fig. 11: Size of 2-core of our networks. The ratio increases with node degree.

C. Related Work

There are many different levels to describe the Internet topology: physical, IP or application. We are interested here in Internet topology at the IP layer. Within this layer, previous works have looked at three different types of topologies: router, AS and PoP.

In this paper, we focussed particularly on PoP-level layer-3 maps. Existing works on PoP maps use traceroute to measure the topology [1]–[5], though they differ in how they infer the topology from the traceroutes. Reverse DNS lookups have been used to group routers into PoPs [1], but this has been shown to be inaccurate [40] as DNS naming is not always linked with geo-location. iPlane [2] groups routers into PoPs by using the TTL values from the routers to measurement vantage points – routers inside a PoP should have similar TTL values to the vantage points. Yoshida *et al.* [4] used delay measurements from vantage points to infer PoP topologies of ISPs in Japan. Most recently, Shavitt *et al.* [5] used a more structural approach to infer PoP maps by first identifying geographically nearby interfaces, then using several heuristics based on the assumption that routers inside a PoP often form a particular structure. Even though the latter approaches of inferring PoP maps are more accurate than the former, they all suffer from measurement errors, biases and inference errors. Our dataset differs from existing work in that we obtain the maps directly from network operators and therefore avoid these errors.

Comparing our results to those of the two previous studies that provide PoP-level graphs across a significant sample of networks [1], [5] we see that our networks are broadly similar in terms of size. The 10 networks in [1] have PoP size varying between 10 and 121, with 6 networks having more than 25 PoPs and 2 more than 100 PoPs. These networks are slightly larger than the networks in our dataset. On the other hand, the networks in our dataset seems to be larger than those in [5], with some cases being much larger.

We have not found any other study of the detailed statistical properties of PoP-level graphs, and there is little point in comparing apples and oranges (or apples and elephants) by comparing our results to router-level or AS-level graphs that are commonly reported in the literature (for example see [1], [16], [21], [41], [42]).

V. CONCLUSION

This paper describes a new data set – the Internet Topology Zoo – based on manual transcription of public network maps. It contains 232 networks at present, and is still growing. We have already used this dataset to perform statistical analysis of PoP-level network topologies.

The Zoo will provide many opportunities for future work including consideration of networks at other levels such as the physical level, and for cross-comparisons between levels.

Most importantly, the Zoo provides a resource for those who wish to validate measurements or test algorithms on real networks.

ACKNOWLEDGEMENT

The authors would like to acknowledge the support of the Australian Research Council through grants DP0985063 and DP110103505, and two Australian Postgraduate Awards.

REFERENCES

- [1] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *ACM SIGCOMM*, (Pittsburg, USA), August 2002.
- [2] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An information plane for distributed services," in *OSDI*, November 2006.
- [3] D. Feldman and Y. Shavitt, "Automatic large scale generation of Internet PoP level maps," in *IEEE Globecom 2008*, pp. 2426–2431, 2008.
- [4] K. Yoshida, Y. Kikuchi, M. Yamamoto, Y. Fujii, K. Nagami, I. Nakagawa, and H. Esaki, "Inferring PoP-level ISP topology through end-to-end delay measurement," in *PAM 2009*, pp. 35–44, 2009.
- [5] Y. Shavitt and N. Zilberman, "A structural approach for PoP geolocation," in *IEEE Infocom 2010*, 2010.
- [6] W. Willinger, D. Alderson, and J. Doyle, "Mathematics and the internet: A source of enormous confusion and great potential," *Notices of the AMS*, vol. 56, no. 5, pp. 586–599, 2009. <http://www.ams.org/notices/200905/rtx0905005586p.pdf>.
- [7] W. Willinger, "The science of complex networks and the Internet: Lies, damned lies, and statistics," Feb 2010. <http://www.maths.adelaide.edu.au/matthew.roughan/workshops.html>.
- [8] "AARNET Annual Report," tech. rep., AARNET, 2009.
- [9] "Research and Education Advanced Network New Zealand Ltd, Annual Report," tech. rep., REANNZ, 2010.
- [10] "BELNET Annual Report," tech. rep., BELNET, 2008.
- [11] CESNET, "CESNET Annual Report," tech. rep., CESNET, 2009.
- [12] "Connecting People, Enabling Innovation, 2009-2010 Annual Report," tech. rep., CANARIE, 2010.
- [13] "ARNES Annual Report," tech. rep., ARNES, 2008.
- [14] E. Gilbert, "Random graphs," *Annals of Mathematical Statistics*, vol. 30, pp. 1441–1444, 1959.
- [15] B.M.Waxman, "Routing of multipoint connections," *IEEE J. Select. Areas Commun.*, vol. 6, no. 9, pp. 1617–1622, 1988.
- [16] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *ACM SIGCOMM*, pp. 251–262, 1999.
- [17] A. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, 1999.
- [18] R. Albert, H.Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *Nature*, vol. 406, pp. 378–382, 2000.
- [19] S.-H. Yook, H.Jeong, and A.-L. Barabási, "Modeling the Internet's large-scale topology," *PNAS*, no. 99, pp. 13382–13386, 2002.
- [20] E. W. Zegura, K. Calvert, and S. Bhattacharjee, "How to model an internetwork," in *Proceedings of IEEE Infocom '96*, (San Francisco, CA), 1996.
- [21] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: degree-based vs. structural," *SIGCOMM Comput. Commun. Rev.*, vol. 32, pp. 147–159, August 2002.
- [22] J. C. Doyle, D. L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, "The 'robust yet fragile' nature of the Internet," *Proceedings of the National Academy of Sciences of the USA (PNAS)*, vol. 102, pp. 14497–502, October 2005.
- [23] V. Cerf and B. Kahn, "Selected ARPANET maps," *Computer Communications Review (CCR)*, vol. 20, pp. 81–110, 1990.
- [24] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "Netscope: Traffic engineering for IP networks," *IEEE Network Magazine*, pp. 11–19, March/April 2000.
- [25] A. Shaikh and A. Greenberg, "OSPF Monitoring: Architecture, Design and Deployment Experience," in *Proc. USENIX Symposium on Networked System Design and Implementation (NSDI)*, March 2004.
- [26] "Internet Protocol." IETF RFC 791, September 1981.
- [27] F. Baker, "Requirements for IP version 4 routers." IETF, Network Working Group, Request for Comments: 1812, July 1993.
- [28] C. Smith, "Traceroute - whitepaper." <http://www.informatik.uni-trier.de/~smith/networks/tspec.html>.
- [29] V. Jacobson, "Traceroute." <ftp://ftp.ee.lbl.gov/traceroute.tar.gz>, 1989-04.
- [30] CAIDA, "Skitter." <http://www.caida.org/tools/measurement/skitter/>.
- [31] yWorks, "yEd Graph Editor." http://www.yworks.com/en/products_yed_about.html, 2010. [Online; accessed 3-Dec-2010].
- [32] M. Himsolt, "GML: A portable Graph File Format," tech. rep., Universitat Passau, 2010.
- [33] U. Brandes, M. Eiglsperger, I. Herman, M. Himsolt, and M. Marshall, "GraphML Progress Report Structural Layer Proposal," in *Graph Drawing*, pp. 109–112, Springer, 2002.
- [34] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring network structure, dynamics, and function using NetworkX," in *Proceedings of the 7th Python in Science Conference (SciPy2008)*, (Pasadena, CA USA), pp. 11–15, Aug. 2008.
- [35] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy, "Revealing the autonomous system taxonomy: The machine learning approach," in *In Passive and Active Measurement (PAM) Workshop*, 2006.
- [36] A. Dhamdhere and C. Dovrolis, "Ten years in the evolution of the Internet ecosystem," in *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement, IMC '08*, (New York, NY, USA), pp. 183–196, ACM, 2008.
- [37] C. Linnaeus, *Systema naturae per regna tria naturae :secundum classes, ordines, genera, species, cum characteribus, differentiis, synonymis, locis*. Holmiae (Laurentii Salvii), 10th edition ed., 1758. (in Latin) Retrieved September 22, 2008.
- [38] A. H. Rasti, N. Magharei, R. Rejaie, and W. Willinger, "Eyeball ASes: From geography to connectivity," in *ACM Sigcomm IMC*, (Melbourne, Australia), 2010.
- [39] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, pp. 167–256, June 2003.
- [40] M. Zhang, Y. Ruan, V. Pai, and J. Rexford, "How DNS misnaming distorts Internet topology mapping," in *USENIX Annual Technical Conference*, May/June 2006.
- [41] P. Mahadevan, D. Krioukov, B. Huffaker, and A. Vahdat, "Orbis: Rescaling degree correlations to generate annotated Internet topologies," *SIGCOMM Comput. Commun. Rev.*, vol. 37, pp. 325–336, January 2007.
- [42] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, k. c. claffy, and A. Vahdat, "The Internet AS-level topology: three data sources and one definitive metric," *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 17–26, January 2006.

PLACE
PHOTO
HERE

Simon Knight received a Bachelor of Engineering (Telecommunications) in 2008, and Bachelor of Economics in 2009, both from the University of Adelaide, Australia. He is currently a PhD Student in the department of Electrical and Electronic Engineering at the University of Adelaide, where he is researching automatic configuration of large-scale emulated networks.

PLACE
PHOTO
HERE

Hung Nguyen is a senior research fellow at the School of Mathematical Sciences, the University of Adelaide, Australia. His research interests include network measurements, tomography, and privacy preserving techniques. He has a PhD in computer science from the Swiss Federal Institute of Technology, Lausanne, Switzerland (EPFL) and is a member of the IEEE.

PLACE
PHOTO
HERE

Dr Nick Falkner is a Lecturer at the University of Adelaide, Australia. His research interests include knowledge representation, network taxonomy, data models, wireless sensor networks, data stream management and semantic annotation for data management. He lectures in a range of courses, including advanced C/C++ programming, computer networking and distributed systems.

PLACE
PHOTO
HERE

Rhys Bowden is a Ph.D. student at the University of Adelaide. His research interests include optimization-based network topologies, network tomography and compressive sensing.

PLACE
PHOTO
HERE

Matthew Roughan joined the School of Mathematical Sciences at the University of Adelaide in 2004. Prior to that he worked at AT&T in the United States. His research interests lie in measurement and modelling of the Internet, and his background is in stochastic modelling with his PhD being in Applied Probability from the University of Adelaide, awarded in 1994.