This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

10 lessons from 10 years of measuring and modeling the internet's autonomous systems

PLEASE CITE THE PUBLISHED VERSION

<http://dx.doi.org/10.1109/JSAC.2011.111006>

PUBLISHER

© IEEE

VERSION

AM (Accepted Manuscript)

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Roughan, Matthew, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 2019. "10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems". figshare.
<https://hdl.handle.net/2134/13142>.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

10 Lessons from 10 Years of Measuring and Modeling the Internet’s Autonomous Systems

Matthew Roughan Walter Willinger Olaf Maennel Debbie Perouli Randy Bush
University of Adelaide AT&T Labs-Research Loughborough University Purdue University IIJ

Abstract—Formally, the Internet inter-domain routing system is a collection of networks, their policies, peering relationships and organizational affiliations, and the addresses they advertise. It also includes components like Internet exchange points. By its very definition, each and every aspect of this system is impacted by BGP, the de-facto standard inter-domain routing protocol.

The element of this inter-domain routing system that has attracted the single-most attention within the research community has been the “inter-domain topology”. Unfortunately, almost from the get go, the vast majority of studies of this topology, from definition, to measurement, to modeling and analysis, have ignored the central role of BGP in this problem. The legacy is a set of specious findings, unsubstantiated claims, and ill-conceived ideas about the Internet as a whole.

By presenting a BGP-focused state-of-the-art treatment of the aspects that are critical for a rigorous study of this inter-domain topology, we de-mythify in this paper many “controversial” observations reported in the existing literature. At the same time, we illustrate the benefits and richness of new scientific approaches to measuring, modeling, and analyzing the inter-domain topology that are faithful to the BGP-specific nature of this problem domain.

Index Terms—Internet topology, modeling, BGP, routing measurements, inference limitations

I. INTRODUCTION

The term “Internet” means (many) different things to (many) different people. Even within the networking community, the term is often used ambiguously, leading to misunderstandings and confusion and creating roadblocks for a genuinely scientific treatment of an engineered system that has revolutionized the way we live.

While mathematics in the form of graph theory has been equally culpable in adopting the use of this vague nomenclature, the “new science of networks” has popularized it to the point where phrases like “topology of the Internet” or “Internet graph” have entered the mainstream science literature, even though they are essentially meaningless without precisely-stated definitions. For one, “Internet topology” could refer to the connectivity structures encountered in any of the seven OSI (Open Systems Interconnection) layers, from the physical fiber and cable connections at the physical layer, all the way to the virtual or logical connections associated with applications such as the WWW (World Wide Web), P2P networks (e.g., BitTorrent), or social media networks (e.g., YouTube, Facebook, Twitter) at the application layer. Moreover, by the very nature of this layered architecture of the Internet, these different connectivity structures are shaped by different sets of technological, economical, and societal forces and evolve in response to different sets of external and

internal signals and responses. Each one of them only offers a (different) glimpse at a global critical infrastructure whose overall purpose and functionalities are determined by a set of layer-specific protocols that run on millions of devices to ensure connectivity among billions of users.

When trying to establish a precise meaning or interpretation of the use of “Internet topology,” in much of the existing literature, we find that the phrase has often been taken to mean a virtual construct or graph created by the Border Gateway Protocol (BGP) routing protocol. Commonly referred to as the inter-domain or Autonomous-System (AS) topology — named after the logical blocks (ASes) that are used in BGP to designate the origin and path of routing announcements — it is this particular connectivity structure that we focus on in this paper.

Our motivation for concentrating on this topology is twofold. First, there has been a wealth of myths, misconceptions, and misinformation in the literature to date about the AS-level Internet. The fundamental problem has been the uncritical reliance on BGP as the main source of measurements. By its very design, BGP is an information-hiding rather than an information-revealing routing protocol, and using it for mapping the Internet inter-domain topology is a “hack” and not a purposefully designed measurement methodology. Being a hack, it should come as no surprise that even though the resulting data contain some amount of useful AS-related information, it lacks critical Internet-wide routing state information necessary for synthesizing the AS-level Internet. Second, when carefully considering the specifics of the AS-level Internet and making proper use of the existing domain knowledge in this area, especially with respect to BGP, this particular topology becomes an amazingly rich source of exciting new problems whose solutions can be expected to provide a deep understanding of critical issues (e.g., resilience, behavior under real-world threats, future evolution) that will be paramount for designing tomorrow’s Internet.

We are not the first to focus on the Internet AS-level topology. In fact, there has been a number of excellent prior reviews that provide a detailed description of the available measurements and of the commonly-used modeling approaches and analysis techniques that have been considered to date for studying the AS topology (see, for example, [1]–[4]). While some overlap with this earlier work is unavoidable, the goal of this review is different. We seek to reflect on the lessons learned from the past 10+ years of Internet topology research. We will be destructive as well as constructive in our criticism of previous efforts. On the destructive side, we aim to debunk aspects of the published research concerning

the AS-level Internet that have created general excitement among networking and non-networking researchers alike, but which fall apart when scrutinized by domain experts, or tested with carefully vetted data. On the constructive side, our effort is aimed at directing and guiding future AS topology research into gainful new areas. To this end, we show how the shortcomings in the existing state-of-the-art have generated openings for a more rigorous, scientific treatment of the AS-level Internet.

A. 10 Lessons from 10 Years of Studying the Internet Autonomous System

Much of the published research on the AS-level Internet seems convincing and sound. After all, the work adheres in large parts to the traditional scientific method; that is, it is typically grounded in real measurement data and follows a generally-accepted modeling approach. However, to a critical networking researcher, it is the very nature of this scientific method which invites questions that probe the soundness of each and every facet of AS-related findings reported in the existing literature. Despite their simplicity, these questions reflect some two decades worth of experience with Internet-related measurement and modeling and can be succinctly summarized as follows. “Are the available measurements appropriate for the purpose at hand?” and “Is modeling or model validation reduced to some basic data-fitting exercise?”

The following is a necessarily subjective list of 10 lessons that we have learned from about a decade worth of published research on the AS-level Internet, with special focus on the data that has fueled much of the research in this area. Our list includes the major topic areas where the AS-level Internet has played a prominent role. The observations on which we base our lessons have become widely accepted within and outside the networking research community, and a main goal of this paper is to “de-mythify” each of these “facts” by examining in detail how they fare in light of the above-stated critical questions.

- 1) The notion of “inter-domain topology of the Internet” is ambiguous, at best, without more precise definitions of terms than typically provided.
- 2) The commonly-used practice of abstracting ASes to generic atomic nodes without any internal structure is an over-simplification that severely limits our ability to capture critical features associated with real-world ASes such as route diversity, policy diversity, or multi-connectivity.
- 3) The traditional approach of modeling the AS-level Internet as a simple connected di-graph is an abstraction incapable of capturing important facets of the rich semantics of real-world inter-AS relationships, including different interconnections for different policies and/or different interconnection points. The implications of such abstractions need to be recognized before attributing network-specific meaning to findings derived from the resulting models.
- 4) The BGP routing data that projects like RouteViews or RIPE RIS have collected and made publicly available are of enormous practical value for network operators, but

were never meant to be used for inferring or mapping the AS-level connectivity of the Internet. The main reason for this is that BGP was not designed with AS-level topology discovery/mapping in mind; instead, BGP’s purpose is to enable ASes to express and realize their routing policies without revealing AS-internal features and, to achieve this goal in a scalable manner, BGP has to hide information that would otherwise aid topology discovery.

- 5) The traceroute data that projects like Ark (CAIDA), DIMES, or iPlane have collected and made publicly available have been a boon to network researchers, but are inherently limited for faithfully inferring or mapping the AS-level connectivity of the Internet. The main reason for this is that traceroute was not designed with Internet topology discovery/mapping in mind; instead, it is a diagnostic tool for tracking the route or path (and measuring transit delays) of one’s packets to some host, and to achieve this diagnostic task, traceroute can ignore issues (e.g., interface aliasing) that would need to be solved first were topology discovery its stated objective.
- 6) Significant additional efforts are required before current models of the Internet’s inter-domain topology derived from the publicly available and widely-used measurement data can purposefully be used to study the performance of new routing protocols and/or perform meaningful simulation studies. At a minimum, such studies need to be accompanied by strong robustness results that demonstrate the insensitivity of reported claims to model variations that attempt to address or remediate some of the known shortcomings of the underlying models or data.
- 7) When examining the vulnerability of the Internet to various types of real-world threats or studying the Internet as a critical infrastructure, it is in general inappropriate to equate the Internet with a measured AS topology. In fact, meaningful investigations of most vulnerability-related aspects of the Internet typically require taking a more holistic approach to Internet connectivity, accounting for details of the physical infrastructure, of how physical connectivity maps to various types of more virtual connectivity, of protocol-specific features, and of traffic-related aspects that manifest themselves at the different connectivity structures.
- 8) While there is a valid role for “observational” studies of the Internet’s Autonomous System, the results of such studies are in general hard to interpret. A more promising method involves performing controlled experiments that allow one to discriminate alternative explanations for results and prevent the effects of one confounding factor from drowning out the effects of others.
- 9) Studies which start with a definite application, and proceed to collect the best data available for that application have shown a much higher rate of success than “fishing expeditions”; that is, studies that target datasets collected by third-parties and analyze them for the sake of analysis.
- 10) In an environment like the Internet where high-variability phenomena are the rule rather than the ex-

ception and where the quality of the data cannot be taken for granted, it is paramount to apply data-analytic methods that have strong robustness properties to the known deficiencies in the observations and naturally account for the presence of extreme values in the data.

The list contains elements that range from definition and meaning of the graph, to measurements and their appropriateness for AS-related studies. They may seem repetitive, but we prefer to err on the side of explicitly stating the problems, rather than leaving the issues implicit.

Debunking past mistakes goes hand in hand with identifying new and interesting directions for more purposeful and promising studies of the AS-level Internet, and our main motivation for this paper is ultimately the latter and not the former. Importantly, we believe that addressing the current deficiencies will move us from treatments of the AS topology as an uninspiring and often meaningless abstract graph towards an approach that views the AS Internet as an economic construct that is constrained by socio-technological factors and is driven by economic incentives and business decisions made by the major players in this area (e.g., service and content providers, large corporations, governments). Although this notion has been advanced by the networking operator community for some time [5]–[9], the networking research community has been slow to react and to distance itself from the popular graph view of the inter-domain topology (examples of exceptions include [10]–[12]).

II. BGP CRITICAL CHARACTERISTICS AND IMPLICATIONS

Routing in the Internet is undertaken on two scales: within an administrative domain or AS and between ASes. Separate routing protocols (and separate routing tables) are used by an AS to spread information about internal and external destinations. The routing protocol used within an AS is termed an Interior Gateway Protocol (IGP) and is the choice of the individual AS. The current de-facto standard inter-domain routing protocol is the Border Gateway Protocol (BGP) [13]–[16], and it is this protocol that concerns us in this section.

A. A Highly Scalable and Expressive Information Hiding Protocol

The characteristics of BGP that have made it successful are at the same time those which significantly add to its complexity. BGP was designed as an “information hiding” protocol, and it is very good at it. The features that contribute to the information hiding capability include: (a) Scalability: BGP runs on a distributed system whose size is the Internet. Each autonomous network that is part of the Internet computes the routes through which it can access the rest of the Internet. To keep BGP scalable, only best paths towards a destination are propagated. (b) Hiding of internal network structures: BGP allows networks to exchange routing information between them without revealing strategic information about their own networks. For example, ASes are often not willing to share sensitive business data, such as the number of routers inside their network or their networks’ topological structure, which customers are buying transit, traffic demands and/or routing

strategies, the location of a company’s data centers and all available paths to reach a particular destination. An example of how deceiving this information hiding capability can be is given in [17]. The authors give examples of what they call “induced updates”, where it is impossible for a monitor to pick-up the root-cause. (c) Configuration flexibility: BGP provides the operators with the expressive power they need to be able to implement the complex and evolving business policies ASes have with each another. By design, BGP hides the information about these policies, and unless this information is published somewhere else and independent from BGP (e.g., on the company website or in the IRR [18]), it is in general difficult or close to impossible to retrieve or infer it. Note that Internet Service Providers (ISPs) today often have a much richer set of policies and value-added services than widely discussed in the literature [6], [19]. Section III-D highlights common assumptions researchers make about business policies.

A BGP-speaking router operates by taking the information about existing routes from its BGP neighbors, the IGP and other sources. The router then makes a decision about which of these provides the “best” route to each destination. These best routes are then (subject to export policies) passed to the router’s BGP neighbors. The process iterates until a stable routing solution is found (if the protocol converges, which is not in fact guaranteed [20], [21]). The output of the BGP’s decision process involves *policies* that can result in routes that may be far from shortest-paths.

The BGP best path selection process [16] and features are well known [13]–[16], and they are critically important when dealing with AS-related inferences. In the following, we list a number of observations that are relevant for the remainder of this paper:

- Every BGP-speaking router in the Internet obtains some information by this routing process and uses this information to route packets. However, this information passed on by BGP is “selective”, not complete, in the sense that a neighboring router only receives the output of a complex selection process and not the various inputs.
- BGP announcements carry useful information. Most notably, for AS topology-related work, they include AS-path information.
- BGP can behave badly, based on non-local policies. For instance, an external policy change over which AS’s network administrator has no control can cause an unexpected and/or undesirable shift in traffic on the AS’s network. The lack of global transparency makes BGP very hard to debug [22].
- The AS-path in the BGP update messages is inserted for loop detection and also to provide some form of distance metric. However, there is no guarantee that traffic is actually flowing along that AS-path [19].

B. BGP Monitor Limitations

As a service to ISPs and the operational community, in the late 1990’s two organizations started collecting and distributing (near) real-time BGP routing information gathered from a number of backbone networks. Those projects are widely referred to as RouteViews [23] and RIPE RIS [24]. They

deployed *BGP Route Monitors* (also called *Collectors*), which are essentially just normal PC's running routing software such as Quagga [25]. Initially, only large providers connected to this service and provided feeds, but later on this service also became popular at IXPs (Internet eXchange Points).

Typically, these monitors record all the BGP update information they receive from their neighbors. They do not announce any prefixes, they do not send or receive traffic—although there are exceptions [26]. The resulting records are the primary source of BGP data that many researchers use.

Both the RouteViews and RIPE RIS projects have collected amazing datasets in the terabytes range providing approximately a decade worth of Internet-wide BGP routing information. The data is currently collected from several hundreds of vantage points in the Internet and is publicly available in open data formats. The data has been a huge boon to network operators debugging network configurations, and the continued reliability of the Internet literally depends on these datasets.

A common characteristic of the RouteViews and RIPE RIS projects is the explicitly- and specifically-stated purpose for collecting BGP routing information in the first place. Both projects were originally motivated by interest on the part of operators in determining how the global routing system viewed their prefixes and/or AS space (see [23], [27]). Importantly, both projects have been silent about the use of their data for mapping the inter-domain topology of the Internet, and for good reasons. First and foremost, the data obtained from a BGP monitor has many limitations, arising principally from the nature of BGP itself, and include:

- 1) The monitor can only see what the connected router chooses to send along. Care is needed when interpreting what is in the data. Contrary to popular belief, one does *not* see the Internet as seen by the connected router. At best, it may be possible to anticipate what a downstream neighbor might receive¹.
- 2) The type of information that can be collected is also not always the same. Most feeds are, what is often called, a “full-feed” (or “default-free” routing table [28]). However, some feeds are “partial feeds” and first go through some filtering process before being sent on to the collector. For example, this may happen at IXPs where the feed is set up for other ASes to show them which routes would be learned if a particular form of peering agreement were signed.
- 3) Some ASes are very large and span multiple continents. Operators often aim at keeping traffic reasonably local, so that the view that is collected in, say, New York, might be very different from a view that is collected within the same AS in, say, Tokyo.
- 4) The current monitors are connected in only a few locations, and each monitor has only a limited viewpoint. Moreover, the locations of these monitors are not randomly distributed across the Internet, but are biased

towards larger *core* networks and IXPs.

- 5) The connections between BGP monitors and routers are not 100% reliable. Session resets, collector down times, and missing updates cause missing data among other problems [29]. Verifying and cross-checking the BGP data is essential, otherwise the data could easily be misinterpreted [30], [31].
- 6) BGP is a path-vector protocol, which means that a single triggering event may cause multiple updates to be observed at a collector, depending on timer-states, topology, and vendor implementation [26], [32].
- 7) Various artifacts appear in the data. With respect to deriving the inter-domain topology from BGP data, one instance is of particular concern: *path poisoning* [33]. Path poisoning is a technique used by some operators and researchers to announce prefixes that contain misleading AS-path information to trigger (false) loop detection at remote ASes.

Generally speaking, the above limitations fall into two categories: artifacts and systematically missing data. Artifacts are not easy to fix, but may with care be handled. On the other hand, systematically missing data is very hard to deal with. In view of the wide-spread use of this data by researchers for the purpose of studying the inter-domain topology of the Internet, we note here that when relying on third-party data that was collected for a specific purpose and using it for a very different purpose, a key question that needs to be asked and answered [34] is “are the existing measurements (which were collected for a specific purpose) of sufficient quality for the purpose for which I want to use them, and how do the defects in the data affect the inferences I intend to make?”

The popular approach of simply using the available RouteViews or RIPE RIS data when trying to study AS-level topology has been pursued for more than a decade without answering the above question. Clearly, taking the data at face value and deriving from them results that are actually trustworthy is highly problematic. The problem is in interpretation, that is converting the data into useful information. This process is fraught because BGP was designed for a particular process (routing) and not for mapping the inter-domain topology. The information contained in the protocol is not the information that most AS topology investigators would ask for. As noted in [35], “what we can measure in an Internet-like environment is typically not the same as what we really want to measure (or what we think we actually measure).” In much of science, we make do with what we can get, but nevertheless, it is important to consider whether such an investigation sheds light on any real problems, or becomes purely an act of sophistry.

C. Other measurements

Looking glass servers are another source of BGP data (with the same limitations), but the information they provide is generally highly constrained in space and time. The Internet Routing Registries (IRRs) provide useful information in general, but are known to contain a significant amount of stale or incomplete data and are therefore typically of limited practical use [36], [37].

Yet another set of data comes from the data plane using traceroute. Although this is a fundamentally different

¹This is a rough approximation, because updates depend on BGP timers that influence what is sent when (which is typically different for each neighbor). In addition, such a downstream router is typically connected to other routers as well, which are not monitored. Therefore, this “input connectivity” of a router is hard to model as well, as it depends on factors such as, type and function of a router, location and AS-related strategies.

measurement technology, it faces a set of similar, and in fact more severe, limitations than BGP measurements. Like the BGP monitors, `traceroute` [38] is a debugging tool that was never intended to measure topology. Its problems are even more extensive than those of BGP, though outside the scope of this paper (see [35], [39] for details). The fact that `traceroute` returns a router path, while BGP returns a path in AS-hops, suggests that these are orthogonal measurements and therefore complementary. In reality, they just overlap the same problem space. An additional difficulty is mapping IP addresses to ASes [40], [41], which may mean that combining these is impractical. At best, there is the problem that the paths seen by the routing protocol and the data plane may just have different, incompatible meanings.

III. FROM ASes TO AS TOPOLOGIES

In the previous section, we emphasized the critical need to understand the basics of BGP when trying to interpret the BGP routing information that has been collected and made publicly available by projects such as RouteViews or RIPE RIS. Here we focus on the use of this information in past studies that reduce the Autonomous System Internet to a simple graph.

A. The Definition of an Autonomous System

One key ingredient of a graph is its nodes. In the context of the AS-level Internet, it is tempting to simply equate a node with an AS, but this begs the question what an AS really is. Most papers loosely define an Autonomous System as a region in the Internet which is under a single administrative control. The term “administrative control” implies a company with commercial interests in the Internet that operates in this space following certain business strategies and targeting specific market niches. However, the reality is more subtle than this.

From a technical perspective, an AS is often viewed in terms of the AS number (ASN) allocated by IANA (Internet Assigned Numbers Authority) or the Regional Internet Registries (RIRs). An ASN is tendered to enable routing using BGP. An extension of this view is to define an AS in terms of its destination prefixes. This is an association that is measurable, say by `traceroutes` or routing monitors. Exceptions to this view include ASes that number internal links from unannounced address space (such as specified in RFC 1918 [42]). Anycast or Multiple Origin AS [43] provide yet another set of counter-examples to a straight-forward mapping between ASN and address space.

Even when it is well defined, the above clashes with the popular view that associates an AS with a set of routers that appear to the outside as if they formed a single coherent system. The administrative view and the logical address-based view of an AS are often inconsistent. An organization may often own a router which has at least one interface IP address belonging to another organization. In fact, many point-to-point IP links occur across a “/30” subnet. When the link joins two networks, this subnet must be allocated from the IP blocks of one or the other connecting network, and so most such connections result in IP addresses from neighboring ASes appearing locally. The problem is exacerbated when an

ISP manages the edge router of one of its customers. Such problems make it exceedingly hard to even define the edge of an AS, let alone measure it.

Another problem arises from the fact that although an AS is often considered to correspond to a single technical administrative domain, i.e., a network run by one organization, it is common practice for a single organization to manage multiple ASes, each with their own ASN [44]. For instance, Verizon Business (formerly known as UUNET) uses ASNs 701, 702, 703 to separate its E-BGP network into three geographic regions, but runs a single IGP instance throughout its whole network. In terms of defining nodes of a graph, these three networks are all under the same operational administrative control, and hence should be viewed as a single node. On the other hand, as far as ASNs are concerned, they are different and should be treated as three separate nodes. The situation is actually more complex since corporations like Verizon Business may own some 200+ ASNs [44] (not all are actually used, though). In many of these cases, a clear boundary between these multiple ASes may not really exist, thus blurring the definition of the meaning of a node in an AS graph. Similar problems can arise when a single AS is managed by multiple administrative authorities which consist of individuals from different corporations. For example, AS 2914 is run partially by NTT/America and partially by NTT/Asia.

All this presumes that an AS is a uniform, contiguous entity, but that is not necessarily true. An AS may very well announce different sets of prefixes at different exit points of its network, or use BGP to balance traffic across overloaded links (other reasons for non-homogeneous configurations are reported in [28]).

For all these reasons, it should be clear that modeling an AS as a single generic node without internal (or external) structure is overly simplistic for most practical problems. Moreover, these issues cannot simply be addressed by moving towards graph representations that can account for some internal node structure (such as in [45]), mainly because BGP is unlikely to reveal sufficient information to infer the internal structure for the purpose of faithful modeling.

B. AS Connectivity

We have seen that the definition of an AS is fraught with problems. Assuming for the time being that the concept of an AS is well defined so that it makes sense to equate each AS with a node in a graph, then what is the set of links? Unfortunately, the question of which ASes are “connected” also has no simple answer, and defining the meaning of a “link” between two ASes requires further consideration.

Does a connection mean the ASes have a business relationship, physical connectivity, connecting BGP session, or that they share traffic? All the above are reasonable definitions, and none are equivalent.

A common construction is an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where the nodes or vertices $\mathcal{V}$ are the ASes and the edges $\mathcal{E}$ are the connections between ASes. Two ASes are said to be connected (at a particular time), if they can exchange routing data (and presumably IP traffic) without the help of an intermediary AS that provides *transit*. In essence, what

we have just defined is a representation of the BGP routing structure as a simple graph. The main question is whether or not this abstraction is of any practical use or relevance.

C. AS Graph Meaning and Extensions

It should be clear from our earlier arguments that the abstraction of the BGP-routing structure of the Internet to a simple graph loses a great deal of information and is an overly simplistic way to view the Internet. All we are seeing is the BGP routing structure of the network, and it is wrong to assume this is somehow the fundamental topology of the Internet. In fact, this graph representation is not a particularly useful topology for any practical purpose. We can make it more purposeful by being more careful about its definition.

First, the AS graph should really be a multigraph. It is very common for two ASes to be connected by multiple links and in different geographic locations [6], [46], [47]. The idea is clearly illustrated by Figure 1 in [46], which shows a “pancake” diagram of the North American Internet backbone. This fact has often been ignored when considering topics such as reliability of the AS graph under link or node failures [48], although it is a necessary ingredient for such studies. Apart from the need to quantify the number of redundant paths available, a failure of a physical link between two ASes might not be visible to any external observer at all (maybe not even throughout the ASes themselves), while a small change in IGP cost could trigger many hot-potato route changes visible in large portion of the Internet [49]–[51]. Perhaps the reason this critical aspect of the topology is typically ignored is that it is hard to measure—BGP monitor data is in general blind to this facet of the topology.

Second, the AS graph should really be a hypergraph. A single “edge” can connect multiple ASes. This is common in an IXP—physical infrastructures managed by third parties where networks can choose to peer with one another for the purpose of exchanging traffic directly, and essentially for free, compared to using some upstream service provider at a cost. At IXPs multiple networks are joined together at one physical location [9], [52]–[54]. One might argue that they are joined by a switch/router, each using point-to-point links, but in at least some cases, that switch has no place in a purely AS-level graph, and so must be considered as a hyperlink between more than two ASes.

Third, ASes are not atomic. It should be clear that an AS is a geographically distributed entity [55], but the problem is even deeper. ASes are comprised of multiple components (routers and the like) distributed over some area of space. In principle (according to the RFCs) the AS should have one routing policy. Not only is it not exactly clear what this means, but it is clearly not true [28], [45], [52]. The components of an AS may not even be contiguous. An AS may rely on a provider AS for transit of its traffic between multiple otherwise disconnected components.

Lastly, there is no clean 1:1 mapping between “network” and “organization” and “AS” [44], [52]. A single organization may use multiple ASes to implement its network, or it may acquire a number of ASes as a result of mergers and acquisitions. A single AS may also represent multiple networks or

companies. It is not uncommon for a small network operator to have a single transit provider, in which case there is often no need for them to route using BGP. In this case, they don’t need an ASN, and they simply appear to be part of the provider AS.

In summary, the AS-graph by itself and as defined above is not particularly useful. It *may* have some scientific interest (though this should be tempered by an understanding of what it really is), but it is not applicable by itself, not to Internet-relevant problems in particular. To be useful, say for predicting network behavior under policy changes or failures, one must also understand something of the policy relationships. It is that point we address next.

D. Policies and Relationships

If the AS-graph by itself and as defined earlier is to be of any use in conjunction with analyzing or controlling BGP routing, we must label the edges and nodes with policies. Although different engineers can define many different policies, these policies must be implemented through BGP, and hence the possible implementations are less rich or varied than the possible policies. It has been common to approximate the range of policies between ASes by a simple set of three relationships: (a) customer-provider, (b) peer-peer, and (c) siblings. This reduction was at least in part motivated by Huston [7], [8] and has been used in various places [9], [56]–[58].

While many relationships fall into these three categories, there are frequent exceptions [45], [52], for instance, in the form of partial transit in a particular region [5], [19]. One way that the partial transit relationship can be implemented is as a hybrid between the customer-provider and the peer-peer relationship: the subscriber receives routes from the provider’s customers and peers, but not the provider’s providers.

Routing under policies (a)–(c) has the “valley-free property”, which both leads to estimation algorithms and simplifies the behavior of BGP. Although the “valley-free property” is taken for granted in a number of papers, there are cases which argue for the opposite². One research study [60] that attempted to “quantitatively characterize BGP announcements that violate the so-called valley-free property” stated that “valley announcements are more pervasive than expected” even in the biased dataset collected from BGP monitors (see Section II-B). They report that 14 out of the 15 ASes that they classify as tier-1 propagate valley announcements. They also recognize that, apart from misconfigurations, there are intentional valley announcements and they attribute them to complex policies of middle-sized intermediate providers.

There are reasons why researchers still use this simplified set of business relationships. Several studies need to model “intent”; for example, in routing security one needs a model of the intended routing to show attacks against it. The simplified model of relationships provides an easy way to achieve such a goal. The critical point when using this model, though, is to show that the results of the study do not rely on it despite using it. For example, a researcher could enhance the

²It was publicly reported [59] that when PSI depeered AboveNet, Verio gave AboveNet transit to/from PSI.

experiments with additional topologies and show that the work is not particularly sensitive to the model assumptions.

Forgetting for the moment the simplification in assuming all policies fit this model, and the simplifications the AS-graph itself makes, the relationships can be represented in the graph by providing simple labels for each edge. In this case, the literature starts categorizing ASes into “tiers” [9], [56]. The concept of “tiers” starts with the tier-1 networks, often defined as those that don’t buy transit from any other AS. As such these networks must peer with each other in a clique. Then, below these are the tier-2 providers who are customers of the tier-1 networks, but peer with each other to some extent, and so on for as many levels as your model suggests.

Another approach is to infer a more generic set of policies consistent with routing observations using a more detailed set of routing measurements [45], [61] and estimate performance by comparing predicted routes to real routes (held back from the inference process).

E. The AS Graph Set

When we consider the above, we start to understand the challenge of giving a precise meaning to the notion of “the” AS graph. In reality, there are multiple incarnations of this graph, each with its own meaning, structure, potential applications, and inference problems. Table I lists some of the possible graphs, all of which have ASes as nodes.

- *business relationship graph*: in its simplest form this graph simply indicates (by an edge) that a business relationship exists between the corporations that own two ASNs. Edges could be usefully labelled by the type of business relationship, and we list a small subset of the possible relationships in Table I.
- *physical link-level graph*: this graph indicates whether two ASNs have a physical (layer 1) connection, and how many such connections they have. The multiple nature of such connections leads this to being a multigraph. The fact that some physical connections are through entities, such as IXPs, that connect multiple ASes leads to this graph being a hypergraph. The graph’s edges could be usefully annotated with link capacity and potentially other features such as geographic location.
- *connectivity graph*: this graph indicates that layer-2 connectivity exists between two ASNs. In many cases the layer-2 connectivity between ASNs would be congruent with the layer-1 connectivity, but with recent advances in network virtualization this may not hold for long [62].
- *BGP routing graph*: the edges in this graph indicate pairs of ASes that have an active BGP session exchanging routing information (i.e., a BGP session that is in the ‘established’ state [16]).
- *policy graph*: the edges in this graph are the same as those in the BGP routing graph, but include directed policy annotations [63]. We define this separately from the BGP routing graph because it may require a multigraph to allow for policy differences between different regions.
- *traffic graph*: it is the same as the BGP routing graph, but the edges are annotated with the amount of traffic exchanged between the corresponding ASes.

Obviously, the definitions above are arbitrary, and could be changed, but they are used here to highlight the ambiguity behind the term “AS graph” and its manifold meanings. Most commonly, what is meant by the AS topology is the structure of the routing graph, possibly with some elements of the policy graph. However, it appears unusual for studies to even define precisely what graph they examine (exceptions being papers such as [9], [64], [65] where the BGP routing graph is explicitly considered).

Rare studies have tried to capture other views, e.g., [47] using other methods such as traceroute. However, as noted in Section II-C, since traceroute has its own set of problems when used to map either intra- or inter-domain topologies, the results of such studies need to be examined very carefully and with full knowledge and explanation of the limits and applicability of the technique.

IV. LESSONS LEARNED ABOUT AS-GRAPH SNAPSHOTS

As noted in the Introduction, there are a number of lessons researchers have learnt over the last decade. Lessons 1-3 have been discussed at length already. Lessons 4-7 are the topic of this Section. The core issue is that BGP was not intended for the purpose of measuring the AS graph. In using data from both RouteViews and RIPE RIS in ways other than intended, the question is raised about its suitability. We have commented extensively on the limitations of those measurements, but here we will examine the impact of these limitations.

As far as we know, the first researchers to use BGP route monitor data for topology-related work were Govindan and Reddy [64]. They introduced the notion of the *inter-domain topology* defined as “the graph of domains and the inter-domain peering relationships.” Specifically, they defined a link in this graph to signify route exchange between the corresponding domains. The paper is quite specific about the nature of the problem that the authors were interested in³. Since the authors hypothesized that two characteristics of the routing system (i.e., the inter-domain topology and route stability) impact Internet wide-area communication, they needed to understand, among other things, features of the Internet’s inter-domain topology and relied on available BGP route monitor data “to derive an *approximate* characterization of the inter-domain topology.”

The paper that coined the term “Internet topology” and is undoubtedly better known and more widely cited than [64] is Faloutsos *et al.* [69]. While this paper is largely responsible for launching much of the subsequent significant research activity in this area, it is also responsible for advancing the alluring notion that the inter-domain topology of the Internet is a well-defined object and can be *accurately* obtained and reconstructed from the available BGP route monitor data. In fact, starting with [69], the *approximate* nature of the inferred inter-domain topology and the limitations of the underlying BGP route monitor data emphasized in [64] have been largely ignored, and the majority of later papers in this area typically only cite [69] and no longer [64]. Unfortunately,

³“A study of the Internet inter-domain routing system; that is, the collection of domains, their policies and peering relationships, and the address prefixes they advertise.”

Graph	Edge Annotation	Graph Type
business relationship	subsidiary, partner, customer,...	directed graph
physical link-level	link capacity	multi- hyper-graph
connectivity graph	-	multigraph
BGP routing graph	-	undirected graph
policy graph	BGP policies	directed multigraph
traffic graph	traffic volumes	directed graph

TABLE I: Example elements of the set of AS graphs.

Paper	Date	Measured	Estimated
Zhang <i>et al.</i> [66]	2004-10-24	45,058	55,388
He <i>et al.</i> [37]	2005-05-12	47,199	59,500
Mühlbauer <i>et al.</i> [45]	2005-11-13	49,241	58,903
Roughan <i>et al.</i> [67]	2004-01	38,397	42,818
	2005-01	45,814	54,582
	2006-01	50,129	59,319
	2007-01	57,038	68,856
	2008-01	63,536	76,944
Dhamdhere <i>et al.</i> [68]	end of 2007	70,000	-

TABLE II: Past estimates of links in the AS-graph.

as commented in [34], the impact of such secondary citations in the measurement arena is especially severe, since critical information available in the original work is often obscured or forgotten. In this context, even a very limited examination of the main Internet topology-related publications in the field of “network science” (e.g., [70], [71]) is illuminating.

The most notable problem in the AS-graph measurement and inference is missing data, primarily missing edges. Most *reachable* ASes appear in a BGP monitor’s view⁴, and so when we combine data from multiple monitors it is unlikely that more than a few active ASes are missing. However, there are a very significant number of missing edges [9], [45], [54], [65]–[67], [73], [74].

Table II shows several estimates reported in these papers and obtained using different techniques, for instance using additional sources of data: IRRs, and Looking Glasses, as well as additional route monitors (some 1,000 in [45]), or statistical techniques [67]. While the approaches used are very different, most of these studies come up with similar numbers on the order of about 20% of the total number of edges, but there is no ground-truth to verify or falsify the accuracy of such estimates. However, if the recent study [54] that focuses exclusively on discovering missing links at IXPs is any indication, the number of missing edges may be significantly larger – quite likely between 50-100%.

We know that the above missing edges cause significant problems in inferring the AS graph. For instance, it is a requirement that a network be multi-homed to obtain an ASN. This means the AS needs to intent to connect to at least two upstream providers. In this sense a “single-homed stub-AS” does not exist. Without any doubt, there are exceptions to this rule. However, the second link is often a backup link which is invisible to BGP outside of the immediate connection, because of BGP’s information hiding⁵. Thus, it may appear as if a large number of ASes are single-homed stubs.

In [65] the authors separate the missing links into ‘hidden’ and ‘invisible’ (for a given set of monitors). The important point about invisible links is that these are links that are missing from the data for structural reasons, that is, it is not just a question of quantity (i.e., numbers of monitors) but quality (i.e., location of monitor). In [67] the authors divide links into a number of classes based on their observability and develop a class based estimator for each of these. We can, thus, place reasonable lower-bounds on the numbers of links that are missing from the data, and know that many such links

exist, but it is hard to put a tight upper bound on the number, because we cannot see what we just cannot see.

The problem with missing links is much more serious than if those links were “missing at random”. In particular, the bias in the type of links that are missing is critical when calculating some metrics on the graph, such as distances, precisely because such links are often *designed* to cut down on the number of ASes traffic must traverse.

Typically, the next step after inferring network topology is to infer policies between ASes. The most common approach to this problem is to assume the universality of the peer-peer, customer-provider, sibling-sibling model, and to infer the policies by finding an allocation of policies consistent with the observed routing [9], [57], [58], [74], [76], [77]. Once relationships are established, a seemingly reasonable next step is to estimate the hierarchical structure as in [9], [56]. However, the effect of large numbers of (biased) missing links has not really been considered in these algorithms. In fact, the tier structure of the Internet seems to be largely an illusion. Recent work has shown that there is little value in the model at present [9], [12], [78]; but, in contrast to the claims of these papers, there is no strong evidence that the situation has actually changed or that the tier model was ever a good model.

The tier-1 ISPs may be a realistic concept, though searches for cliques amongst groups of large providers always produce smaller sets than any reasonable grouping suggested by the nature and scope of the companies involved. In reality, even if the tier-1 concept is correct, it ignores the transitory nature of the network and the business relationships that need to be maintained to keep connections alive.

There is also a natural set of “bottom tier” ISPs who do not provide transit to any other BGP speaking ISP (they certainly can provide transit to other ISPs, just not the ones that appear as separate ASes). We sometimes call these “stub” ASes (though note that even in the simple AS-graph they may not be degree one nodes). These actually form the vast majority, some 30,000 of the 36,000 or so ASes do not appear to provide transit today.

However, it is very difficult to classify the intermediate transit-providing ASes. Certainly there are ambiguities because one AS may appear to be in different tiers based on its relationship with various other providers, and there can be no consistent labeling as a result. More serious though, are the problems with the whole model that assumes that all relationships are of these types—as we have noted ASes are often not homogenous.

Other analyses of the AS-graph have included studies of its reliability [48]. Once again, ignorance of the approximations

⁴This is because typically an AS announces at least one prefix. There are of course exceptions, as discussed in Section III-A. Furthermore, note that it is not guaranteed that all ASes are reachable from all other ASes [72].

⁵Note that complex BGP policies may play a role in this as well [22], [75].

in the AS-graph (for instance of the multilink nature of the real connections), and the problems in measurements (the number of missing edges) invalidate such studies completely. Likewise for other graph-based metrics [79] applied without understanding the above issues.

V. LOOKING AHEAD

Given our list of problems described here, one might be tempted to think that the AS-graph and routing data in general are useless until these datasets are drastically improved. However, apart from their operational utility, RouteViews and RIPE RIS have provided the essential ingredients for many important studies that match the services' goals [80]. A number of these studies have improved the Internet significantly, and in the majority of such successful papers there is no need to exploit the "graph" view of the network. Examples include: (a) The discovery of slow convergence and persistence oscillation in routing protocols [9], [20], [81]–[86]. (b) Understanding of the impacts (positive and negative) of route flap dampening [87], [88]. (c) Determining how much address space and how many ASNs are being actively used [89]. (d) Looking for routing "Bogons" often related to Internet address hijacking [90]–[94]. (e) Debugging network problems [9], [17], [95], [96].

On the measurement side, there have also been many advancements towards improving our view of AS topology. For instance:

- 1) As BGP routing changes, often multiple potential paths are explored and these paths (which are unlikely to actually be used as a final choice) can show some of the alternative routes available in the network [66], and thus a more complete topology.

There is an unfortunate side-effect of this type of measurement. It introduces a Heisenberg-like uncertainty principle. It is not clear whether observed changes are due to the micro-phenomenon of path exploration, or macro-phenomena of link changes, new entrants, etc. The longer we make observations, the more complete they may seem, but we then do not know whether all of those links existed at the same time. Such uncertainty principles appear to be present in a number of Internet measurement contexts [97] where we trade off "accuracy" of the measurements against "time localization". This approach does not overcome the structural bias.

- 2) Missing edges can be found using additional datasets, e.g., RIRs and looking glasses [37], [66], [73], [98], or IXP data [9], [37], [54], [98], though care must be exercised with any additional dataset.
- 3) Beacons [26], [28], [82]: a routing beacon is just a router that advertises and withdraws certain prefixes on a regular schedule. Examination of the observed announcements and withdrawals by various route monitors then allows estimates of protocol behavior such as convergence time.
- 4) Route poisoning prevents announcement from reaching certain parts of the Internet. As with beacons, it allows one to examine the behavior of BGP in a more controlled manner. This is perhaps the only way to see (some) backup paths, or to understand whether an ISP uses default routing [28], [33].

- 5) There are also attempts to not just estimate the topology but derive some quality measure for the resultant AS-graph [67], [99], [100].

On the one hand, these and other advances on the measurement side suggest that the missing link problem may be solved in the not-too-distant future, paving the way for highly predictable future research efforts focusing on a new round of characterizing and modeling these "more complete" AS graphs. However, when trying to understand the reasons for the various advancements on the measurement front and examining the sources that yield the improved data, it becomes increasingly obvious that genuine advances on the research front will not come from "more of the same"—traditional studies of the Internet's AS graph as a graph-theoretical construct devoid of most features or attributes that make it relevant and interesting from a networking perspective. Instead, the latest measurement efforts and resulting data all highlight the fact that the AS-level Internet is much richer and rewarding than what can be described with a simple di-graph. Providing a mathematical framework that fully reflects and respects that richness and supports the search for the main technological and economic factors that shape the AS-level Internet and are responsible for its evolution will be at the heart of new scientific advances in this area. The reward of these new efforts promises to be a unique ability to successfully reverse-engineer this critical Internet construct for the purpose of strategically influencing its future functioning and evolution.

In addition to defining a rather unconventional agenda for future research in this area, the recent advancements on the measurement side listed above also relate directly to lessons 8-10. First, controlled experiments (i.e., experiments that have a "control" sample against which the experimental data can be compared) are necessary in order to precisely derive which factors of interest affect which variables. Controls allow one to discriminate alternative explanations for results, and prevent the affects of one confounding factor drowning out the affects of others (see [26], [28]). This is basic tenet of the scientific method, but seems to have been ignored in this area of research. Most studies have been "observational", and while there is a valid role for such experiments, for instance in epidemiology, they are intrinsically harder to interpret.

Second, we have much more hope for studies that set out to measure a particular phenomena, or solve a particular problem (for some instances see [9], [101]), and which design their measurements around that problem than we do for "fishing expeditions" which simply take a set of data, and mess around with it until they find something apparently of interest. The latter approach is more often uncritical of the flaws in the data, because at the end of the day the results are often treated uncritically, whereas results aimed at solving a particular problem are assessed by whether they really solve that problem.

Third, in a world where high-variability phenomena are the rule rather than the exception and where the quality of the data cannot be taken for granted, it is paramount to apply data-analytic methods that have strong robustness properties to the known deficiencies in the observations and naturally account for the presence of extreme values in the data. A common approach in, for instance, machine learning towards providing

better quality measures of performance is to hold out some set of data for testing the quality of inferences made with a set of training data, and this approach has much to be recommended in this context too, despite the fact the idea has been used sparingly (see for instance [45]).

VI. CONCLUSION

The underlying story of this paper may seem to be one of woe and tragedy. We start with flaky data, progress to uninformed, uncritical interpretation, and finish without application. Can anything be done?

The answer is yes. However, in writing this paper we hope to illuminate the critical aspects of the AS-graph that every researcher working in this area or using the available data should know. While past research has identified important and difficult problems, it is the way these problems have been “solved” that we critique in this paper. By emphasizing that constructs such as the inter-domain topology of the Internet cannot be treated justly as simple abstract graphs –devoid of essentially all network-specific meaning– we outline some directions forward towards solving a set of more challenging, interesting, and ultimately more rewarding problems. The lessons of this paper will hopefully form a checklist for any student or new researcher in this area that will enable them to avoid the pitfalls which have reduced the value of some past research. Simply stated, to ensure value of future research in this area, any work on the structure and evolution of the Internet’s Autonomous System has to account for the economic, technological, and social forces that shape this critical element of the Internet.

REFERENCES

- [1] H. Chang, *Modeling Internet’s Inter-Domain Topology and Traffic Demand Based on Internet Business Characterization*. PhD thesis, University of Michigan, 2006.
- [2] B. Donnet and T. Friedmann, “Internet Topology Discovery: A Survey,” *IEEE Communications Surveys & Tutorials*, vol. 9, pp. 56 – 69, 2007.
- [3] A. Dhamdhere, *Understanding the Evolution of the AS-level Internet Ecosystem*. PhD thesis, Georgia Institute of Technology, 2008.
- [4] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, “Network topologies: inference, modeling and generation,” *IEEE Communications Surveys*, vol. 10, no. 2, 2008.
- [5] W. B. Norton, “Internet Service Providers and Peering,” 2010. <http://drpeering.net/white-papers/Internet-Service-Providers-And-Peering.html>.
- [6] W. B. Norton, “A Study of 28 Peering Policies.” <http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html>.
- [7] G. Huston, “Peering and Settlements - Part I,” *The Internet Protocol Journal*, vol. 2, March 1999.
- [8] G. Huston, “Peering and Settlements - Part II,” *The Internet Protocol Journal*, vol. 2, June 1999.
- [9] J. Peha, “Retransmission Mechanisms and Self-Similar Traffic Models,” in *IEEE/ACM/SCS Communication Networks and Distributed Systems Modeling and Simulation Conference*, January 1997.
- [10] H. Chang, S. Jamin, and W. Willinger, “To peer or not to peer: Modeling the evolution of the Internet’s AS-level topology,” in *Proceedings of IEEE INFOCOM*, 2006.
- [11] A. Dhamdhere and C. Dovrolis, “The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh,” in *6th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, (Philadelphia PA, USA), December 2010.
- [12] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, “Internet Inter-Domain Traffic,” in *ACM SIGCOMM*, 2010.
- [13] T. Griffin, “An Introduction to Interdomain Routing and BGP,” in *Tutorial in ACM SIGCOMM*, 2001.
- [14] O. Bonaventure, “Interdomain routing with BGP4,” Advanced one-day course, <http://www.info.ucl.ac.be/people/OBO/BGP/>, Louvain-la-Neuve, May 2003.
- [15] J. Stewart III, *BGP4: Inter-domain Routing in the Internet*. Addison-Wesley, Boston, 1999.
- [16] Y. Rekhter and T. Li, “A Border Gateway Protocol (BGP-4),” RFC 4271, Jan 2006.
- [17] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, “Locating Internet Routing Instabilities,” in *ACM SIGCOMM*, 2004.
- [18] Internet Routing Registry. <http://www.irr.net/>.
- [19] M. Yoshinobu, “What makes our policy messy,” 2010. <http://www.attn.jp/maz/p/c/bgpworkshop200904/bgpworkshop-policy.pdf>.
- [20] T. G. Griffin and G. Wilfong, “An Analysis of the MED Oscillation Problem in BGP,” in *ICNP*, 2002.
- [21] T. G. Griffin and G. Wilfong, “On the correctness of IBGP configuration,” *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 32, no. 4, pp. 17–29, 2002.
- [22] T. G. Griffin and G. Huston, “BGP Wedgies,” RFC 4264, 2005.
- [23] University of Oregon RouteViews project. <http://www.routeviews.org/>.
- [24] RIPE’s Routing Information Service. <http://ris.ripe.net/>.
- [25] “Quagga Routing Suite,” <http://www.quagga.net>.
- [26] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan, “BGP Beacons,” in *ACM SIGCOMM Internet Measurement Conference*, (Miami Beach, Florida, USA), October 2003.
- [27] A. Antony and H. Uijterwaal, “Routing Information Service – R.I.S. Design Note,” <ftp://ftp.ripe.net/ripe/docs/ripe-200.pdf>.
- [28] R. Bush, O. Maennel, M. Roughan, and S. Uhlig, “Internet optometry: assessing the broken glasses in Internet reachability,” in *ACM SIGCOMM Internet Measurement Conference*, (New York, NY, USA), pp. 242–253, ACM, 2009.
- [29] A. Flavel, O. Maennel, B. Chiera, M. Roughan, and N. Bean, “Clean-BGP: Verifying the Consistency of BGP Data,” in *Proc. Internet Network Management Workshop*, 2008.
- [30] J. Cowie and A. Ogielski, “Global Routing Instabilities During Code Red II and Nimda Worm Propagation,” *NANOG 23*, October 2001.
- [31] B. Zhang, V. Kambhampati, M. Lad, D. Massey, and L. Zhang, “Identifying BGP routing table transfers,” in *ACM SIGCOMM Mining the Network Data (MineNet) Workshop*, 2005.
- [32] T. G. Griffin and B. J. Premore, “An Experimental Analysis of BGP Convergence Time,” in *Proceedings of ICNP*, 2001.
- [33] L. Colitti, *Internet Topology Discovery Using Active Probing*. PhD thesis, University di “Roma Tre”, 2006.
- [34] B. Krishnamurthy, W. Willinger, P. Gill, and M. Arlitt, “A Socratic Method for Validation of Measurement-based Networking Research,” *Computer Communications*, vol. 34, pp. 43–53, 2011.
- [35] 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>.
- [36] 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.
- [37] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, “A systematic framework for unearthing the missing links: Measurements and Impact,” in *USENIX/SIGCOMM NSDI*, April 2007.
- [38] V. Jacobson, “Traceroute.” <ftp://ftp.ee.lbl.gov/traceroute.tar.gz>, 1989-04.
- [39] 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>.
- [40] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, H. Zhang, and L. Zhang, “A Framework to Quantify the Pitfalls of Traceroute in AS-level Topology Measurement,” *IEEE JSAC, Special Issue on “Measurement of Internet Topologies”*, vol. 29, this issue, 2011.
- [41] Z. M. Mao, J. Rexford, J. Wang, and R. Katz, “Towards an Accurate AS-Level Traceroute Tool,” in *Proceedings of ACM SIGCOMM*, (Karlsruhe, Germany), August 2003.
- [42] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. d. Groot, and E. Lear, “Address Allocation for Private Internets,” 1996.
- [43] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, “An analysis of BGP multiple origin AS (MOAS) conflicts,” in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, IMW ’01, pp. 31–35, ACM, 2001.
- [44] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger, “Towards an AS-to-Organization Map,” in *ACM Sigcomm IMC*, (Melbourne, Australia), 2010.

- [45] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-topology model that captures route diversity," in *ACM SIGCOMM*, (Pisa, Italy), 2006.
- [46] M. Liljenstam, J. Liu, and D. Nicol, "Development of an Internet Backbone Topology for Large-Scale Network Simulations," in *Proc. of 2003 Winter Simulation Conference*, (New Orleans, LA), 2003.
- [47] N. Spring, R. Mahajan, and T. Anderson, "Quantifying the causes of path inflation," in *ACM SIGCOMM*, (Karlsruhe, Germany), 2003.
- [48] R. Albert, H. Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *Nature*, vol. 406, pp. 378–382, 2000.
- [49] R. Teixeira, A. Shaikh, T. G. Griffin, and J. Rexford, "Dynamics of Hot-Potato Routing in IP Networks," 2004.
- [50] R. Teixeira, A. Shaikh, T. G. Griffin, and G. M. Voelker, "Network sensitivity to hot-potato disruptions," 2004.
- [51] R. Teixeira, N. G. Duffield, J. Rexford, and M. Roughan, "Traffic Matrix Reloaded: Impact of Routing Changes," in *Proc. Passive and Active Measurement Workshop (PAM)*, 2005.
- [52] Y. Hyun, A. Broido, and K. C. Claffy, "Traceroute and BGP AS path incongruities," tech. rep., UCSD CAIDA, 2003. <http://www.caida.org/publications/papers/2003/ASP/>.
- [53] K. Xu, Z. Duan, Z.-L. Zhang, and J. Chandrashekar, "On Properties of Internet Exchange Points and Their Impact on AS Topology and Relationship," *Networking, Springer-Verlag, LNCS*, vol. 3042, pp. 284–295, 2004.
- [54] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?," in *SIGCOMM IMC'09*, pp. 336–349, 2009.
- [55] A. H. Rasti, N. Magharei, R. Rejaie, and W. Willinger, "Eyeball ASes: From Geography to Connectivity," in *ACM Sigcomm IMC*, (Melbourne, Australia), 2010.
- [56] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Infocom*, 2002.
- [57] F. Wang and L. Gao, "On Inferring and Characterizing Internet Routing Policies," in *ACM SIGCOMM/USENIX Internet Measurement Conference*, (Miami, Florida, USA), October 2003.
- [58] J. Xia and L. Gao, "On the evaluation of AS relationship inferences," in *Globecom*, 2004.
- [59] "Nanog Mailing List Archives," November 2007. <http://seclists.org/nanog/2000/Nov/179>.
- [60] S. Y. Qiu, P. D. McDaniel, and F. Monrose, "Toward valley-free inter-domain routing," in *IEEE ICC*, 2007.
- [61] Z. Mao, L. Qiu, J. Wang, and Y. Zhang, "On AS-level path inference," in *ACM SIGMETRICS*, 2005.
- [62] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, "Virtual routers on the move: live router migration as a network-management primitive," in *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, SIGCOMM '08, (New York, NY, USA), pp. 231–242, ACM, 2008.
- [63] T. G. Griffin, "The Stratified Shortest-Paths Problem (Invited Paper)," in *COMSNETS*, (Bangalore, India), January 2010.
- [64] R. Govindan and A. Reddy, "An Analysis of Internet Inter-Domain Topology and Route Stability," in *IEEE INFOCOM*, pp. 850–857, 1997.
- [65] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)completeness of the observed Internet AS-level structure," *Transactions on Networking*, vol. 18, no. 1, pp. 109–122, 2010.
- [66] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level Topology," *ACM SIGCOMM Computer Communication Review (CCR) special issue on Internet Vital Statistics*, January 2005.
- [67] M. Roughan, J. Tuke, and O. Maennel, "Bigfoot, Sasquatch, the Yeti and other missing links: what we don't know about the AS graph," in *ACM SIGCOMM Internet Measurement Conference*, (Vouliagmeni, Greece), October 2008.
- [68] 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.
- [69] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-Law Relationships of the Internet Topology," in *ACM SIGCOMM*, 1999.
- [70] A. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, 1999.
- [71] S.-H. Yook, H. Jeong, and A.-L. Barabási, "Modeling the Internet's large-scale topology," *PNAS*, no. 99, pp. 13382–13386, 2002.
- [72] C. Labovitz and A. Ahuja, "Shining Light on Dark Internet Address Space," *NANOG 23*, 2001.
- [73] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," *Computer Networks*, vol. 44, no. 6, pp. 737–755, 2004.
- [74] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. C. Claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 37, no. 1, pp. 29–40, 2007.
- [75] Cisco Systems, "BGP Cost Community," 2005. http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/s_bgpcc.pdf.
- [76] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *Global Telecommunications Internet Mini-Conference*, 2000.
- [77] G. Battista, M. Patrignani, and M. Pizzonia, "Computing the types of the relationships between autonomous systems," in *IEEE INFOCOM*, 2003.
- [78] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?," in *Proc. PAM'08*, pp. 1–10, 2008.
- [79] M. E. J. Newman, "The Structure and Function of Complex Networks," *SIAM Review*, vol. 45, pp. 167–256, June 2003.
- [80] Goals of the Routing Information Service. <http://www.ripe.net/ripe/docs/ripe-200>.
- [81] K. Varadhan, R. Govindan, and D. Estrin, "Persistent route oscillations in inter-domain routing," tech. rep., 96-631, USC/ISI, 1996.
- [82] C. Labovitz, R. Malan, and F. Jahanian, "Internet Routing Stability," in *Proceedings of ACM SIGCOMM*, 1997.
- [83] C. Labovitz, R. Malan, and F. Jahanian, "Origins of Internet Routing Instability," in *Proceedings of INFOCOM*, 1999.
- [84] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Wide-Area Network Failures," in *Proceedings of FTCS*, 1999.
- [85] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet Routing Convergence," in *Proceedings of ACM SIGCOMM*, 2000.
- [86] K. Varadhan, R. Govindan, and D. Estrin, "Persistent Route Oscillations in Inter-Domain Routing," *Computer Networks*, March 2000.
- [87] Z. Mao, R. Govindan, G. Varghese, and R. Katz, "Route Flap Dampening Exacerbates Internet Routing Convergence," in *ACM SIGCOMM*, 2002.
- [88] C. Pelsser, O. Maennel, P. Mohapatra, R. Bush, and K. Patel, "Route Flap Damping Made Useful," in *Proc. PAM'11*, 2011.
- [89] G. Huston, "IPv4 Address Report," 2007. <http://www.potaroo.net/tools/ipv4/index.html>.
- [90] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *ACM SIGCOMM'06*, pp. 291–302, 2006.
- [91] The Team Cymru, "The Team Cymru Bogon Reference Page." <http://www.team-cymru.org/Services/Bogons/>.
- [92] P. Boothe, J. Hiebert, and R. Bush, "How Prevalent is Prefix Hijacking on the Internet?," *NANOG 36*, February 2006.
- [93] N. Feamster, J. Jung, and H. Balakrishnan, "An empirical study of bogon route advertisements," *ACM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 63–70, 2005.
- [94] GIZMODO: China's Internet Hijacking Uncovered. <http://gizmodo.com/5692217/chinas-secret-internet-hijacking-uncovered>.
- [95] R. Bush, J. Hiebert, O. Maennel, M. Roughan, and S. Uhlig, "Testing the reachability of (new) address space," in *INM'07: Proceedings of the 2007 SIGCOMM workshop on Internet network management*, (New York, NY, USA), pp. 236–241, ACM, 2007.
- [96] M. Roughan, T. Griffin, M. Mao, A. Greenberg, and B. Freeman, "IP Forwarding Anomalies and Improving their Detection Using Multiple Data Sources," in *ACM SIGCOMM Workshop on Network Troubleshooting*, (Portland, OR, USA), pp. 307–312, September 2004.
- [97] M. Roughan, "Fundamental Bounds on the Accuracy of Network Performance Measurements," in *ACM SIGMETRICS*, (Banff, Canada), pp. 253–264, June 2005.
- [98] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "Lord of the links: A framework for discovering missing links in the Internet topology," *IEEE/ACM Transactions on Networking*, vol. 17, no. 2, pp. 391–404, 2009.
- [99] "Internet AS-level topology construction & analysis." <http://topology.neclab.eu/>.
- [100] R. Winter, "Modeling the Internet Routing Topology with a Known Degree of Accuracy - in less than 24h," in *ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS)*, 2009.
- [101] J. H. Park, D. Jen, M. Lad, S. Amante, D. McPherson, and L. Zhang, "Investigating occurrence of duplicate updates in BGP announcements," in *Passive and Active Measurement Conference*, (Zurich, Switzerland), 2010.

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.

Olaf Maennel is a lecturer at Loughborough University in the United Kingdom since September 2009. Before that he was at Deutsche Telekom Laboratories and at the School of Mathematical Science at the University of Adelaide in South Australia. He got his Ph.D. (Dr. rer. net) from the Technical University in Munich (Germany) in 2005. His research interests are routing, active measurements, next generation internet technology, as well as configuration management.

Walter Willinger received the Diplom (Dipl. Math.) from the ETH Zurich, Switzerland, and the M.S. and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, NY. He is currently a member of the Information and Software Systems Research Center at AT&T Labs-Research, Florham Park, NJ. Before that, he was a Member of Technical Staff at Bellcore Applied Research (1986-1996). He is a Fellow of ACM, IEEE, SIAM, and AT&T, and for his work on the self-similar ("fractal") nature of Internet traffic, he received the 1994 W.R. Bennett

Prize Paper Award from the IEEE Communications Society, the 1996 IEEE W.R.G. Baker Prize Award from the IEEE Board of Directors, and the 2005 "Test-of-Time" Paper Award from ACM SIGCOMM.

Debbie Perouli is a Ph.D. student in the Computer Science department at Purdue University, West Lafayette, IN. She received her Bachelor's degree (5 year Diploma) in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), Greece, in 2006. Her research interests include BGP routing, path algebras, network modelling and configuration.

Randy Bush is a Research Fellow and network operator at Internet Initiative Japan, Japan's first commercial ISP. He specializes in network measurement especially routing, network security, routing protocols, and IPv6 deployment. He has been heavily involved in transferring Internet technologies to developing economies for over 20 years.