ResearchGate

See discussions, stats, and author profiles for this publication at: https://www.researchgate.net/publication/290995765

Inferring Business Relationships in the Internet Backbone

ARTICLE *in* INTERNATIONAL JOURNAL OF NETWORKING AND VIRTUAL ORGANISATIONS · FEBRUARY 2016

Impact Factor: 0.5

2 AUTHORS, INCLUDING:

Benjamin Fabian

Humboldt-Universität zu Berlin

80 PUBLICATIONS **493** CITATIONS

SEE PROFILE

All in-text references underlined in blue are linked to publications on ResearchGate, letting you access and read them immediately.

Available from: Benjamin Fabian
Retrieved on: 08 March 2016

Inferring Business Relationships in the Internet Backbone

Martin Huth[1] and Benjamin Fabian[1,*]

[1] *Institute of Information Systems, Humboldt University Berlin, Spandauer Str. 1, 10178 Berlin, Germany.* [*] *Corresponding author; email: bfabian@wiwi.hu-berlin.de*

Abstract

Several researchers during the last decade have encountered the problem of how to infer business relationships between Autonomous Systems (ASes) of the Internet. Since the Internet has a decentralized structure and public data sources containing inter-domain routing information have not been created for topology inference, there are no accurate and comprehensive maps of the Internet readily accessible. This challenge has inspired several approaches for inferring business relationships between ASes from BGP routing data. This article presents one implementation of the most recent and most promising approach for relationship inference on AS-level. The algorithm used has been improved in terms performance and quality of the sanitizing process. Unlike recent projects, not a only snapshot of the topology of the Internet has been inferred but a comprehensive map showing the Internet over the last decade. The correctness of this implementation and the inferred data set is examined by comparison with a business relationship graph and a validation data set provided by related work.

Keywords: Internet Topology, Business Relationship, Autonomous Systems, Internet Measurement

1. Introduction

The Internet has become an indispensable foundation for today's business and society. Not only has it changed the business model of formerly existing companies but also built the basis for establishing completely new business models. Some of the biggest and most valuable companies started as pure Internet firms (e.g., Facebook or Google) [1]. The Internet, as a global critical infrastructure, needs to be studied with respect to robustness and resilience to disasters and attacks [2], power structure [3], information flow, censorship, and topological connectivity of individual businesses [4] or

services such as cloud computing [5]. All these investigations need a solid understanding and mapping of the Internet topology as a vast and complex routing but also business infrastructure [6]. By convention the Internet architecture is often studied on two abstraction levels [7]. On one level, there are so called domains of which each usually hosts a number of routers. The internal routing environment of those domains is organized by an internal routing protocol, the so called Interior Gateway Protocol (IGP). The second abstraction level of this routing structure describes the inter-domain routing environment. Each domain is an autonomous unit in the overall routing context and is therefore called *Autonomous System (AS)*, referred to by a uniquely identifying number – the Autonomous System Number (ASN). This second abstraction level – the AS-level – describes how the local routing domains are connected with each other. It is organized by the Border Gateway Protocol (BGP) [8].

Due to the decentralized organization of the Internet's architecture, there is no accurate map of it accessible. Different research projects have made efforts to map the topology of the Internet. A representation at the AS-level in form of a graph where nodes are ASes and edges are routing links between ASes has been the preferred choice. The main sources for inferring those maps are BGP routing data and traceroute-based measurements [9]. However, this data was never primarily meant to be used for inferring an AS-level mapping of the Internet. The design of BGP was never driven by the mapping idea. Instead it is supposed to serve as a scalable framework to enable communication between AS domains without revealing domain-internal information [10]. Thus, the inference of Internet topologies at the AS-level is somewhat challenging given these inherent limitations.

Moreover, the vast amounts of existing ASes are operated by different administrative entities such as Internet Service Providers (ISPs), companies and universities. These parties have different (often confidential) commercial interests which, in turn, influence and determine the AS-level topology and traffic patterns. Participating ASes engage in certain business relationships with each other. Business relationships between two ASes can be broadly categorized into two main types: customer-to-provider (c2p) and peer-to-peer (p2p) [11]. For more than a decade several research projects have focused on inferring business relationships between ASes. The result of these efforts are different algorithms for business relationship inference on AS-level. Some projects have made their data publicly available (e.g., CAIDA and UCLA). However, the available business-annotated or directed AS-level graphs are usually significantly smaller than undirected AS-level graphs in terms of the number of ASes and links [12, 13].

The goal of this article is to infer a comprehensive directed graph of the AS-level Internet. This topology shall map business relations between ASes accurately. For this purpose the most promising approach for inferring AS-level business relationships is presented and implemented. The implemented script is applied on a collection of BGP data from a time range of nearly 13 years. Subsequently, the inferred graph is validated and examined for its correctness by drawing a comparison to a data set provided by UCLA [13]. Moreover, an evaluation is presented by using a validation data set by [14].

This article is structured as follows. The first section introduced motivation and research goal. The second part deals with the tools and metrics used for this article. Subsequently, background and theoretical concepts in the context of policy-based routing are explained in detail. This is followed by a summary of related work in the field of business relationship inference of the AS-level Internet. The fourth section focuses on the implementation of the selected approach. Details about the data sources and the implemented algorithm are presented. Then, the results acquired using the implemented script are examined. A validation is presented at the end of this part. The final section includes a discussion of the results including advantages and limitations of the used approach.

2. Preliminaries: Tools and Methods

In this section the tools and methodologies are provided. At first, an insight into the used tools for gathering, inferring as well as analyzing the data is presented. Afterwards, metrics important for evaluating the inference algorithm are presented.

2.1. Perl Programming Language

This article mainly concentrates on the implementation of an algorithm for inferring business relationships between ASes. Any implementations for this article have been made using Perl. Perl is an acronym for Practical Extraction and Report Language. Perl is a simple language that has been designed for powerful text manipulation and has been developed by Larry Wall [15]. Despite Perl being considered a scripting language, it shares many characteristics with system programming languages (e.g., C), like the handling of various I/O aspects and process control [16].

2.2. Zebra Dump Parser

The raw source files available or BGP paths are stored in a binary format. For later processing, the files needed to be parsed in order to extract the

Figure 1: Transit degree computed by using AS paths. The node degrees of A and B are both 3. B has a transit degree of 3. A has a transit degree of 2 because it has not been seen announcing D's prefixes to any neighbors. Nodes with a transit degree of zero are stubs [14].

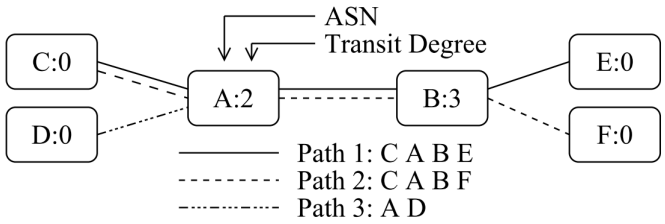

AS paths. For this purpose, the freely available script *zebra-dump-parser.pl* written by Marco d'Itri has been adopted. The script is written in Perl and is capable of parsing the Routing Information Base (RIB) files, like the files provided by RouteViews [17]. The script processes each AS path of each route. The last and second-last AS in the paths are recorded as the origin of the route and its "neighbor" networks (prepends are ignored). The output is a list of AS paths.

2.3. Graph Metrics

For the implemented algorithm, two basic graph metrics were significant: node degree and transit degree. The *node degree* is the total number of neighbors an AS has. The *transit degree*, on the other hand, is the number of unique neighbors that appear on either side of an AS in adjacent links.

Figure 1 illustrates how Luckie et al. compute the transit degree of an AS. The transit degree is used to create a sorted list of ASes for later inference. ASes with a transit degree of zero are called stubs and are considered to build the lowest level in the AS hierarchy. The transit degree is a more accurate metric for the rank of an AS than the node degree because it reduces ordering errors caused by stub networks with a large peering visibility. A BGP peer that shows a view over the Internet from its perspective is called *vantage point* (VP). These stubs provide a vantage point or peer with many vantage points.

2.4. Confusion Matrix Evaluation

In order to evaluate the predictive quality of the implemented algorithm, a confusion matrix can be used. Figure 2 shows the four possible outcomes of a two-case classification scenario. True Positives and True Negatives represent correct classifications. In contrast, False Positives occur when

Figure 2: Confusion Matrix: It shows the possible outcomes of a two-class prediction. True Positives and True Negatives represent correct predictions. False Positives and False Negatives are erroneous predictions (see [18]).

		Predicted Class	
		Yes	No
Actual Class	Yes	True Positive (TP)	False Negative (FN)
	No	False Positive (FP)	True Negative (TN)

incorrectly class "Yes" was predicted but the actual class is "No". False Negatives occur when class "No" was predicted but the true class is "Yes". For validation of the inferred results two metrics were used: True Positive Rate (TPR or Recall) and Predictive Positive Rate (PPV or Precision):

- $Recall = TPR = \frac{TruePositives}{NumberofActualPositives} = \frac{TP}{TP+FN}$
- $Precision = PPV = \frac{TruePositives}{NumberofPredictedPositives} = \frac{TP}{TP+FP}$.

TPR describes the proportion of actual positive values that have been correctly predicted as positive. The PPV, on the other hand, is the proportion of values predicted as positive that are actually positive [18].

In the context of business relationship inference of links between ASes, each inferred type of relationship can be considered as a two-case scenario (e.g., a Confusion Matrix for p2p inferences).

3. Background and Related Work

3.1. Border Gateway Protocol (BGP)

A number of ISPs administrate the Internet where each operates a certain subset of all ASes. Inter-domain communication (i.e., between ASes) is performed by the so-called Border Gateway Protocol (BGP). It is built to exchange routing and reachability information between BGP-"speaking" systems [8]. BGP is specified in several Request For Comments (RFCs) and its current version 4 (BGP-4) is described in RFC 4721 [19].

BGP is designed to enable systems to exchange network-reachability information. This information consists of lists of AS paths and can be translated into a graph of ASes. By that, routing loops can be prohibited and policy-based routing can be enforced. Initially, BGP-speaking systems

Figure 3: Inter-AS business relationship types. Short arrows show which information is shared between ASes (based on [21]).

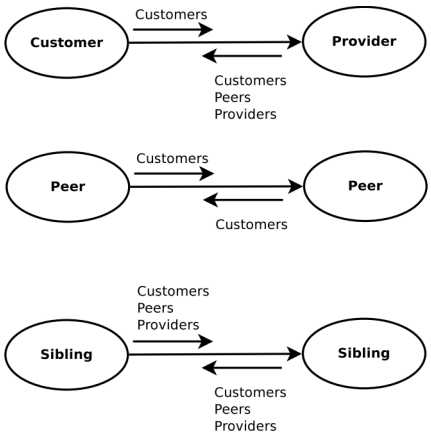

exchange all routing data. After that only incremental updates are sent. These updates are not scheduled regularly, and only the optimal routing path to a network is advertised. Hence, if paths of a router change, the router only advertises the portion of paths that has changed. The metric BGP uses to determine the best path bases on an arbitrary unit number which specifies the degree of preference of a particular link. This degree of preference is assigned to each link by the network administrator. It can be influenced by different aspects like stability, speed, delay, or costs [20].

3.2. Policy-based Routing

On an abstract level the Internet consists of ASes and links between these ASes. However, connectivity between two ASes does not imply reachability itself. A vast number of ISPs is operating individual subsets of all participating ASes [11]. Each ISP engages in different formal and informal relationships with other ISPs. These relationships are usually enforced by certain business agreements which in turn translate into certain routing constraints. Each of these routing constraints influences the traffic flow [12]. Accordingly, business relationships can be mapped on the links between ASes. Gao [11] set the groundwork with a business relationship model that categorizes the types of links between ASes into three distinct categories: *customer-to-provider* (c2p), *peer-to-peer* (p2p) and *sibling-to-sibling* (s2s). Figure 3 illustrates the three types of business relationships between ASes.

The first type c2p is the most common one. A link between two ASes is called c2p if one AS (the customer) is paying the other AS (the provider) for

Figure 4: AS paths with business relationships. (see [11])

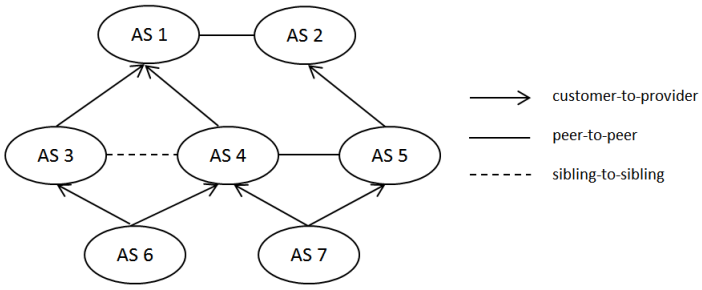

its transit services in order to obtain reachability to ASes the provider can reach directly and to those the provider can reach through its own providers. To achieve this, the provider reveals all his known paths to the customer. The customer will only announce routes of his own customers to its provider. Routes to other peers and providers, on the other hand, stay unknown to its provider.

Another important business relationship between two ASes is referred to as peer-to-peer (p2p). A p2p relationships exists between two ASes if both ASes obtain access to each others' customers and their customers. However, peer and provider routes are confidential and are not revealed to the peer. Typically, neither of them pays the other for transit services. This is advantageous for both peers since they would have to pay a provider for data transit otherwise. Accordingly, this so-called settlement-free peering is attractive for both parties as neither could convince the other to become a customer. The third category, sibling-to-sibling (s2s), is rather rare. It describes the situation when two ASes administratively belong to the same administrative unit. The announcement of routes between the two ASes are not restricted. This is normally the case if two ASes belong to the same ISP [11, 12].

Figure 4 shows examples for all three relationship types within an AS graph. There exists a c2p relationship between AS5 and AS2. Hence, AS5 will offer its route to AS2 and thereby reachability to AS1 since AS2 has a p2p relationship to AS1. Furthermore, AS5 reveals to AS7 the routes it knows through its p2p relationship with AS4. Contrarily, AS5 will reveal its routes to AS2, however, it will detain all routes it learned through its p2p relationship with AS4. AS3 and AS4 will typically share all their routes with each other because of their s2s relationship.

Figure 5: AS paths with business relationships. Paths (3, 6, 4) and (7, 5, 4, 1) are violating the valley-free principal (see [12]).

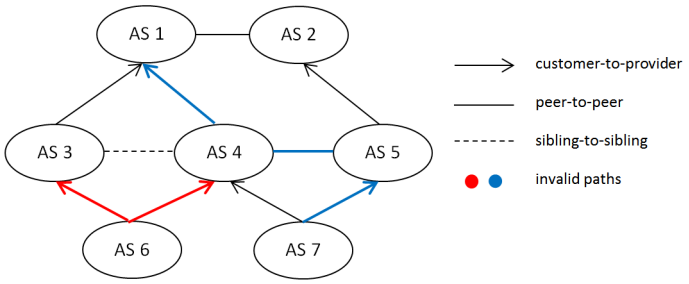

3.3. The Valley-free Model

Based on the aforementioned parameters, Gao [11] defined a valley-free model of the Internet topology on AS-level. Basically, the valley-free model assumes there exist only valid paths within the topology. A valid path is a path where for each AS providing transit services there is at least one AS paying for its services. In a more formal way this can be captured by two requirements: Firstly, a p2c link can only be followed by a p2c link or a s2s link. Secondly, after a p2p there can only appear p2c or s2s links [12, 11].

Figure 5 illustrates examples for both cases of invalid constellations. The red path (3, 6, 4) represents the first case. It is invalid because AS6 would transport traffic for AS3 and AS4. However, neither of them is paying AS6. The blue path (7, 5, 4, 1), on the other hand, violates the second requirement for a valley-free graph. None of the involved ASes is paying for AS4's transit services. Therefore, this path is invalid. The valley-free model is a common assumption used for most of the published approaches for business relationship inference on the AS-level (see Section 3.7).

3.4. BGP Data Sources

The two main projects for collecting BGP data are Route Views [17] and RIPE-RIS [22]. They host so-called *collectors* which collect BGP data by establishing peer-to-peer sessions with operational routers within a significant amount of ASes [9]. Two kinds of files are provided by each collector: RIB files, which are snapshots of the current RIB, and updates. RIBs are usually collected in certain intervals, i.e., every two hours for Route Views [17] and every eight hours for RIPE-RIS [22]. Updates are written in shorter intervals between 5-15 minutes. The file names contain the file type (rib or update), the date and time they were written, e.g., `rib.20040517.1359.bz2`.

Figure 6: Zebra dump parser output variants.

verbose	AS path	AS origin
TIME: 2004-3-2 01:22:07	3.0.0.0/8 2497 7018 80	3.0.0.0/8 80
TYPE: MSG_TABLE_DUMP/AFI_IP	3.0.0.0/8 4777 2497 7018 80	3.0.0.0/8 80
VIEW: 0 SEQUENCE: 0	4.0.0.0/8 2497 3356	4.0.0.0/8 3356
PREFIX: 3.0.0.0/8	4.0.0.0/8 4777 2497 3356	4.0.0.0/8 3356
ORIGINATED: Wed Feb 25	4.17.225.0/24 2497 701 11853 6496 6496 6496	4.17.225.0/24 6496
23:18:48 2004	4.17.225.0/24 4777 2516 3561 11853 6496 6496	4.17.225.0/24 6496
FROM: 202.249.2.169 AS2497	4.17.226.0/23 2497 701 11853 6496 6496 6496	4.17.226.0/23 6496
AS_PATH: 2497 7018 80	4.17.226.0/23 4777 2516 3561 11853 6496 6496	4.17.226.0/23 6496
NEXT_HOP: 202.249.2.169		4.17.251.0/24 6496

The RIBs are stored in the Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format which is described in RFC 6396 [23]. MRT is a binary format and can be translated into ASCII. Route Views [17] and RIPE-RIS [22] suggest three different tools for that: *libbgpdump* – a library written in C, *PyBGPdump* – a Python library, and *zebra-dump-parser* – a script written in Perl (see Section 2.2) [17, 22].

In scope of this article, the zebra-dump-parser has been used. By setting the variable `$format` the output type can be chosen. Three different output types are possible: *verbose* which basically causes an output of all information contained in the original MRT-file, *AS path* which produces a list of AS paths and *AS origin* which produces a list of the mere origins of the paths. Figure 6 shows the three output types. For the conducted inference the second variant has been used. Furthermore, there exists a variable `$ignore_v6_routes` which if set causes ignoring IPv6 (Internet Protocol version 6) routes.

A BGP peer that shows a view on the Internet from its own perspective is called *vantage point* (VP). The first AS in an AS path (see second column in Figure 6) shows the VP of the path.

3.5. Clique (Tier-1 ASes)

As of October 2013 there are more than 70,000 ASNs allocated by IANA to the different RIRs. The sheer size of this number makes it impossible to establish connections to the entirety of all ASes and gather accurate information from all of them. A common assumption has been been prevalent for quite some time that the Internet can be structured into several tiers. Roughan et al. [10] state that this is a misconception. However, Oliveira et al. [9] state that there exists a group of ASes on top of the graph, the tier-1 ASes. They showed that, what they call the *public view*, can be seen fairly completely by those tier-1 ASes over time [9]. Luckie et al. [14] define

the *clique* of the AS-level topology as multiple large transit providers forming a peering mesh. Thereby, customers and indirect customers of a transit provider can reach global connectivity without entering multiple transit provider relationships.

3.6. Customer Cone

Luckie et al. [14] describe a view on the relationship graph that is called the *customer cone*. The customer cone of an AS is the set of all ASes that can be reached by following its customer links, its customers' customer links, and so on. This structure is important for relationship inference in order to mitigate cycles of p2c links. Moreover, it can be used for further analysis of the inferred graph of business relationships. For instance, the customer cone of an AS is a more precise representation of an AS's size than just its degree. It reflects the number of ASes that directly or indirectly pay an AS for its transit services. The bigger an AS's customer cone, the more important its role in the Internet. Accordingly, ASes in the Clique have the biggest customer cones. ASes at the bottom of the hierarchy have no customer cone which reflects the fact that they have to pay for all of their transit.

3.7. Related Work on Business Relationship Inference

For more than the last decade, research has concentrated on the inference of business relationships at the AS-level. This section gives an overview of relevant work in this field. In 2001 Gao [11] was one of the first to set pioneering groundwork for AS relationship inference. She introduced a solution which classified three types of business relationships between ASes: customer-to-provider (c2p), peer-to-peer (p2p) and sibling-to-sibling (s2s). This solution is based on the assumption that BGP paths are hierarchically described by the valley-free model. This model assumes that each path consists of an uphill and a downhill segment. The uphill segment consists of zero or more c2p links or sibling links. On top there are zero or one p2p links followed by the downhill segment consisting of zero or more p2c or s2s links (see Section 3.3). The valley-free model represents typical commercial relationships that exists in the Internet. For each AS providing transit services, there is at least one AS paying for its services. In other words, there is no AS in any AS path that remains unpaid for its transit service.

Gao [11]'s approach intents to maximize the total number of such valley-free paths. This is done by assuming the AS with the largest node degree within a path as the top AS. Furthermore, the algorithm assumes that neighboring ASes within an AS path with a similar node degree (i.e., difference R which is to be fine-tuned) have a p2p relationship. A s2s relationship

between two ASes of a link is inferred if both ASes provide transit services for each other. Gao [11] conducted minor validation by using internal information obtained from the tier-1 AS of AT&T. Furthermore, for validation of the inferred s2s relationships she used information from the ARIN WHOIS service [11].

Xia and Gao [24] published an improved approach based on Gao [11]'s algorithm. They based their inference on partial ground truth information based on commercial agreements between ISPs. The information was obtained from the BGP community attribute, the instances from AS-SET object (both to be found in verbose BGP RIB-files) and the Internet Routing Registry (IRR). This basis served as starting point for the inference process that was conducted using the valley-free model and Gao [11]'s algorithm. They validated 6.3% of their inferences and concluded with an accuracy of their approach of 96.1% for p2c inferences and 89.33% for p2p inferences [24].

Subramanian et al.[25] mathematically formulated the relationship inference as a combinatorial optimization problem referred to as *Type of Relationship (ToR)* problem: Given an AS graph derived from BGP routing data, assign the relationship type (c2p or p2p, ignoring s2s) to each link such that the total number of valley-free paths is maximized. They speculated that this ToR optimization problem is NP-complete.[1] Therefore, they introduced a heuristic-based solution in order to rank each AS. The ranking is based on an AS's apparent distance to the core of the graph assessed from different VPs. Similar to Gao[11]'s approach, adjacent ASes are compared. If two ASes have a similar rank, they are assigned a p2p relationship. Otherwise, a p2c relationship is inferred. Subramanian et al. [25] did not validate their results.

Di Battista et al. [27] later proved that the ToR optimization problem is in fact NP-complete. They presented mathematical approximation solutions to the problem. Furthermore, they showed that inference of p2p relationships under the ToR framework is infeasible. Their solutions only infer c2p relationships and neglect p2p as well as s2s relationship inference. This task was left open to future work. Di Battista et al. [27] also provided no validation of their results.

Dimitropoulos et al. [28] stated that ToR has certain limitations: Ignoring s2s relationships causes proliferation of erroneous inferences. A solu-

[1] NP stands for non-deterministic polynomial time. A problem is NP-complete if it is not solvable in polynomial time in any known way [26].

tion maximizing the number of inferred-to-be-valid paths is not necessarily correct. Furthermore, in situations when several solutions with the same amount of valid paths exist, ToR has no means to deterministically select the most realistic solution. Dimitropoulos et al. [28] approached this problem by manually establishing a dictionary to map syntactically different organization names that belong to one single organization. Then they assigned edges as s2s of ASes that belong to the same organization according to the dictionary. Organizational information was gathered from the IRR. They state that this data is not always up-to-date but reasonably accurate since these information change less frequently than BGP routing information. Thereafter, all s2s edges were removed. Furthermore, Dimitropoulos et al. [28] achieved improvement of c2p inference integrity by enhancing the ToR maximization objective by introducing a second objective. Edges gather a bonus b_i if edge i is directed from an AS with smaller degree to an AS with higher degree. The bonus equals the weight function $f(d_i^-, d_i^+)$ which is a function of the degrees of adjacent ASes to edge i. Value f is high if the difference between d_i^- and d_i^+ is high. Otherwise, it is small and, accordingly, the value of b_i is small. The objectives of the optimization problem are: (O_1) Maximize the number of valid paths in P and (O_2) maximize the sum $\sum_{i \in E} b_i$ (i.e., the number of c2p inferences where the node degree of the provider is larger than the customer).

Dimitropoulos et al. [28]'s approach combined the strength of Subramanian et al. [25], Di Battista et al. [27] and Gao [11] since it uses the valley free model, it attempts to maximize the number of valid paths, and uses the knowledge of AS degree gradients to assign edge directions. Moreover, a Boolean variable for assigning whether a link's direction has to be changed or remains the same was introduced. Furthermore, the multi-objective optimization problem is simplified to the well-known MAX-2-SAT. MAX-2-SAT had been used the first time for AS relationship inference by Wang et al. [29]. MAX-2-SAT is a special type of boolean algebra problem, i.e., 2-satisfiability: the problem of determining whether a collection of two-valued (Boolean or binary) variables with constraints on pairs of variables can be assigned values satisfying all the constraints. They provided validation of 3,724 AS relationships and concluded correct inference of their algorithm for 96.5% of c2p links, 82.8% of p2p links and 90.3% of sibling links. The validation represented 9.7% of the public view and was the most comprehensive validation at that time. The disadvantage of this approach is that MAX-2-SAT's complexity is NP-hard and, thus, inference for recent graphs does not complete in practical time durations [28, 14, 29].

Dimitropoulos et al. [28]'s approach had formerly been in use by CAIDA [12] in order to provide freely accessible AS relationship data. It has been replaced by Luckie et al. [14]'s approach which, also will be used in the scope of this article. It will be presented in more detail in the following section (see Section 4).

Apart from CAIDA, another project that periodically provides AS relationship is UCLA. Zhang et al. [30] describe the method that is used by UCLA for automatically deriving an undirected AS graph from BGP data, i.e., an AS-level graph that does not contain business relationships between ASes. Oliveira et al. [9]'s paper, on the other hand, describes the inference of business relationships at the AS level. Oliveira et al. [9] conjecture that there's a group of ASes on the top of the AS-level hierarchy – tier-1 ASes. The common conception is that these tier-1 ASes do not pay for each other's transit services. Hence, they build a peering clique. A clique of a graph is a subset of the graph that forms a complete subgraph, i.e., each node of the clique has a direct link to every other member of the clique [31]. Oliveira et al. [9] assume that the tier-1 ASes are publicly available. Their algorithm starts from this clique and infers all links that are observable from the clique to be p2c. The remaining links are inferred to be p2p [9]. However, since the number of region-specific c2p relationships that can only be seen below regional ISPs is increasing, Oliveira et al.'s algorithm infers too many p2p relationships [14].

Gregori et al. [21] proposed an approach similar to the one of Oliveira et al. [9]. Briefly, it uses the life time of a path as a metric to infer the business relationships between ASes. The algorithm identifies all possible business relationships and uses the time parameter to actually infer the applicable relationship [21]. Neither of the two latter articles provided validation of the proposed approaches [21, 9].

4. Relationship Inference

This section elaborates on the actual relationship inference. At first, light will be shed on the data sources used. Furthermore, one of the latest algorithms for business relationship inference on AS-level of the Internet will be presented. The algorithm in use was provided by Luckie et al. [14] in a manuscript of the finally published article.

4.1. Data Sources

In Section 3.4 the main sources for BGP data, RouteViews and RIPE-RIS, were described. In order to restrict runtime of the scripts only Route-

Views has been used for gathering BGP data. At the time of data collection, RouteViews provided BGP data from 13 different collectors. The RIB files can be accessed either via FTP or HTTP access. For this implementation the HTTP archive has been used.

The files contain a time stamp representing their creation date. However, hour and minute when the files were written vary from hour to hour, day to day and collector to collector. Therefore, the script first assembled a list of file names from every 2nd day of each month of every year from every collector. Luckie et al. [14] used a time window of the first five days of every collector for one month to extract a set of AS paths. They used the 5-day window in order to see even unstable paths which represent backup paths. For this article only one file per month instead five were download because of the fairly big number of files that had to be processed. In total, 972 RIB files were downloaded by the script. The average size of each compressed file was $\approx 16 MByte$ or $\approx 300 MByte$ uncompressed. It turned out that even by limiting to only one file per month the derived AS set grew fairly big due to the fact that a time range of almost 13 years has been considered.

Furthermore, the Internet Assigned Numbers Authority (IANA) provides an up-to-date list of ASN assignments of the available ASN pool. It provides information about which number blocks are assigned to which RIR [32]. Within the scope of this article the list was used for sanitizing of AS paths.

Moreover, a list of known Internet Exchange Points (IXPs) [33] in combination with an AS name list [34] is used to assemble a comprehensive list of IXPs. Like IANA's list this set of ASes serves the sanitizing process.

4.2. Algorithm and Implementation

As mentioned earlier, several approaches base their inference algorithms on the valley-free concept initially introduced by Gao [11]. In contrast, Luckie et al. [14] developed a new method for business relationship inference. Instead of the valley-free model their method is based on three assumptions:

1. There is a clique of large transit providers at the top of the AS-level hierarchy. These large transit providers build a p2p mesh so that their customers and indirect customers can reach global connectivity without having multiple transit providers.

2. Most customers enter into a transit agreement to be globally reachable. Besides clique ASes, every AS requires a transit provider to reach global connectivity. Accordingly, if an AS A becomes a customer of

provider B, the provider announces routes to A to all its providers or to its peers if B is part of the clique. However, backup and region-specific relationships are exceptions to this assumption.

3. Cycles of p2c links (e.g., an AS graph consisting of A, B, C where A is customer of B, B is customer of C and C is customer of A) do not exist [14].

Luckie et al.'s approach [14] does not try to maximize valley-free (hierarchical) paths. Therefore, inference can be conducted using path triplets (e.g., $A - B - C - D \rightarrow (A - B - C), (B - C - D)$). Triplets contain all necessary constraint information, and processing of triplets is less complex than processing of whole AS paths considering CPU time and memory usage [14].

4.2.1. Sanitizing

Before BGP paths can be processed, they have to be sanitized in certain aspects. Luckie et al. [14] divide the sanitizing of paths in four parts. At first, they discard paths containing artifacts such as:

- Loops indicating path poisoning, i.e., in a path (A-**B**-C-**B**-D), AS B would prevent the path from being chosen by a non-adjacent upstream AS C since B appears twice in the path.
- ASes that are reserved for private use and ASes within unallocated ASN space should not appear in AS paths. [32]'s list [32] is used as a reference to discard such paths.
- IXPs are considered as artifact ASes since the actual relationships exist between the ASes participating in the exchange. Luckie et al. [14] manually assembled a list of 25 IXP ASes. IXP ASes from this list are removed from valid paths.

Moreover, removal of path padding is applied for paths containing the same ASN twice, e.g., $(A - \mathbf{B} - \mathbf{B} - C - D) \rightarrow (A - \mathbf{B} - C - D)$.

In a second step, the algorithm creates a sorted list of all ASes. Order criteria are transit degree, node degree, and ASN to break ties. Thirdly, they infer the clique based on calculating the maximal clique using the Bron/Kerbosch algorithm [35]. The relationship between each two members of the clique is set to p2p. Finally, after inferring the clique, paths are

discarded where any two clique members in a path are separated by a non-clique AS. This condition indicates a poisoned path since ASes in the clique by definition do not pay for transit.

For our article, Luckie et al.'s approach [14] has been adjusted in several aspects. In addition to the aforementioned artifacts, multiple-origin ASes (MOAS) have been considered as such. For a path starting with a MOAS set, clear identification of the announcing BGP peer is ambiguous [36]. Multi-origin paths have the form ({A,B}, C, D). In addition, all ASes in ASDOT format were converted to ASPLAIN format. ASPLAIN and ASDOT are two different textual representations of ASNs. ASPLAIN is a syntax scheme representing ASes using decimal integer notation (e.g., ASN of value 66847 would be represented by the string "66847"). In contrast, ASDOT uses two integer values joined by a period character. The integer value on the left is a multiplier of 2^{16}. The integer on the right side of the period character represents a decimal value. The value of the right integer added the product of multiplier and 2^{16} returns the ASPLAIN format (e.g., ASPLAIN representation 1.1311 is converted by $1 * 2^{16} + 1311 = 66847$).

Instead of manually assembling a list of IXPs to exclude from AS paths, a list of known IXPs by the research project of Augustin et al. [37] has been used. Their data is publicly available ([33]). The list does not contain the ASNs of the IXPs. However, Potaroo [34] provides a list of AS names. By mapping the IXPs' names to their ASNs, a list of 63 ASNs could be derived and used for the sanitizing process.

Finally, the 21 clique ASes provided in the manuscript of Luckie et al. [14] have been used as the clique. This is tenable because of the small number of clique ASes and the stable appearance of these ASes in the clique over time.

4.2.2. Inference

This subsection presents the actual inference steps. *Notation:* In the following, a p2p relationship between two ASes X and Y will be represented by $X - Y$, a p2c relationship by $X > Y$ and a c2p relationship by $X < Y$. In the relationship $X > Y$ the provider will be represented by X. Y represents the customer, accordingly. A link with no inferred relationship will be referred to by X ? Y.

Using the set of triplets, the actual relationship inference is conducted in seven steps which will be presented in detail in the following. One important data structure is the so called customer cone (see Section 3.4). The customer cone of an AS is the transitive closure of its customer relations, i.e., the set of all of its customers, and its customers' customers, and so on. The customer

Figure 7: Example for inferring providers, customers, and peers from sanitized BGP paths. The sorted order of ASes is assembled following transit degree, node degree, and ASN (adapted from [14]).

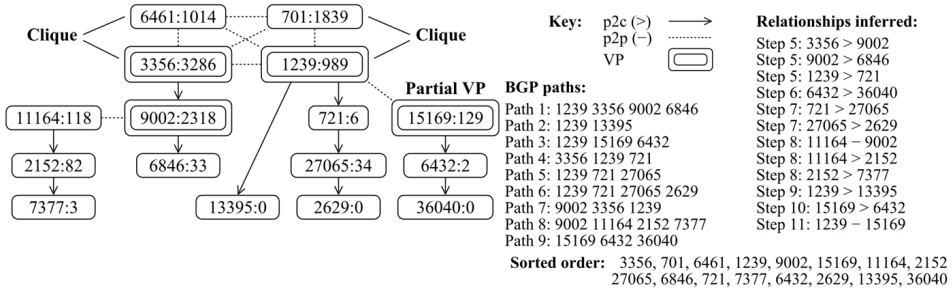

cone is important when it comes to preventing cycles of p2c links: Whenever a c2p between two ASes is inferred, the customer and all members of the customers' customer cone are added to the provider's customer cone. A c2p relationship will not be inferred if the provider is already in the customer cone of the potential customer to avoid a p2c cycle.

In order to improve runtime and memory usage of this algorithm, an inverted representation of this structure has been implemented for this article, i.e., a *provider cone*. The provider cone of an AS represents the set of its providers and its providers' providers. Accordingly, the condition testing is done the other way around, i.e., a c2p will not be inferred if the potential customer is already in the provider cone of the potential provider. Hence, the provider cone delivers all necessary information. Access time for membership testing is linear with moderate memory usage. Only using a customer cone would mean a trade-off between access time and memory usage.

For the remaining part of this section, Figure 7 is used to illustrate each step of inference.

Step 1 − c2p relationships top-down: . In the first step of the algorithm, each AS Z of the sorted list of ASes will be visited top down (i.e., ordered by transit degree, node degree and ASN to break ties). Clique members will be skipped since they have by definition no providers. The algorithm infers $Y > Z$ if a triplet of the form $X - Y\ ?\ Z$ or $X > Y\ ?\ Z$ exists. The former is the case at this point if X and Y are both clique members. The latter can be observed if Y has been visited in a previous iteration of this step. Visiting ASes by descending rank is important in order to avoid false inferences due to router misconfigurations in one direction of a p2p relationship. This could

be the case if an AS leaks provider or peer routes to a peer. In that case, a c2p relationship will be inferred if the provider or peer is observed to be closer than the customer to at least one VP in at least one triplet. This method is based on the assumption that each AS enters a relationship with a provider to achieve global connectivity, i.e., at least one VP should observe the provider announcing the customer's routes.

Step 1 infers 90% of all c2p relationships [14]. This step, at first, considers AS 9002 since it is the first non-clique AS in the sorted list. The first triplet of path 1 is used for inference as it matches the pattern X − Y ? Z (see Figure 7).

Step 2 − c2p relationships from VPs announcing no provider routes:. In the second step, a c2p relationship is inferred from VPs not announcing any provider. This step is based on either of the following two assumptions: (1) Partial VPs export only customer routes (i.e., the vantage has a p2p session with the collector) or (2) the VP has a default route to a provider. Hence, it exports customer and peer routes to the collector (i.e., p2c relationship with a collector). The algorithm infers Y > Z for each path X ? Y ? Z where X is a partial VP[2] and Z is stub (i.e., the transit degree of Z is 0).

Path 9 and AS 15169 being a partial VP leads to the inference 6432 > 36040 in this step (Figure 7).

Step 3 − c2p relationships to smaller degree providers:. In the third step, c2p relationships are inferred for links in the very unlikely case that a provider has a smaller transit degree than the potential customer. If a triplet W > X ? Y is observed where the transit degree of X is smaller than the transit degree of Y and the triplet (W, X, Y) is a suffix (i.e., at least one path ends with (W, X, Y)), this step infers X > Y. In this case, also, Y > Z will be inferred where X > Y ? Z is observed [14]. Step 3 infers 721 > 27065 using path 5 and 27065 > 2629 using path 6 (Figure 7).

Step 4 − Customers for provider-less ASes:. Step one and three require a provider for c2p relationship inference. Step four assigns relationships to links with provider-less ASes. Provider-less ASes can occur within some regional and research networks. At first, the provider cone is used for quickly assembling the subset of ASes which do not have a provider, yet. The algorithm visits each neighbor W of each provider-less AS X top-down (along descending rank of W). If a triplet (W, X, Y) is observed, the algorithm

[2]A partial VP provides routes to fewer than 2.5% of all ASes [14].

infers W − X and X > Y because W has not been observed announcing X to providers or peers before. In this case, furthermore, Y > Z will be inferred where X > Y > Z is observed [14]. The provider-less AS 11164 in Figure 7 and path 8 lead to the inference of 9002 − 11164 as well as 11164 > 2152. Further, 2152 > 7377 is inferred.

Step 5 − c2p relationships for stub-clique links:. Stub ASes are extremely unlikely to achieve a clique AS's peering requirements. Therefore, ASes with direct links to a clique member are inferred to be a customer. In this step, X > Y will be inferred for an unknown relationship where X is a clique member and Y is a stub. Step 5 uses path 2 to infer 1239 > 13395 (see Figure 7).

Step 6 − Collapse adjacent links with no relationships:. Step six looks for triplets with unresolved links X ? Y ? Z and attempts to infer Y > Z. For each triplet X ? Y < P where P is different from Z, the algorithm infers X < Y. If such a triplet X ? Y < P does not exist and there exists a triplet Q ? Y ? X, no inference is made because this would imply Y > X and the original triplet would resolve to X < Y and Y > Z. However, only one side of the original triplet can be resolved with confidence. If neither X ? Y < P nor Q ? Y ? X can be observed, the algorithm infers X > Y [14].

In this step the relationship 15169 > 6432 is resolved (Figure 7).

Step 7 − p2p relationships for all other links:. The last step simply infers p2p for all remaining unresolved links. Looking at the example in Figure 7, this step infers 1239 > 15169.

5. Results

In this section, the results of the relationship inference are elaborated. Initially, data provided by UCLA was used as a reference, later also validation data provided by Luckie et al. [14]. In the following, characteristics of the inferred data set will be illustrated. At first, a comparison regarding relationships between the two data sets will be drawn. Secondly, in order to get an insight into the differences and similarities, single AS nodes will be examined. UCLA provides a monthly updated data set of observed AS links and inferred business relationships on their website (UCLA [13]). According to Luckie et al. [14], this algorithm performed closest to their algorithm during validation. Although UCLA provides monthly dumps, the data only grew in the number of new links tagged as "unknown" during our data collection.

Finally, a comparison of our data with the validation data set provided by Luckie et al. [14] will be presented.

5.1. AS Business Relations

The outcome of our parsing, sanitizing and inferring of 972 MRT files from 13 Route Views archives (2001-2013) is a total of 300,903 unique links. 224,577 of them were inferred to be c2p and 76,326 p2p relationships. This leads to a ratio of 74,63% of c2p relationships and 25,37% p2p relationships.

As a reference for comparison, the UCLA data set of August 2013 was used. UCLA's data is stored as a table of three columns. Column one and two contain ASNs where each row represents a link between two ASes. The third column contains the labels for the inferred business relationship, i.e., c2p, p2c, p2p or unknown. Each link in UCLA's data set is stored in a forward and backward representation. In an AS graph, the link between two ASes A and B is represented by one edge which has a certain direction. However, the relationship type can be read by either using A or B as a starting point, i.e., A is a customer of B or B is a provider of A, respectively. Accordingly, the link would be stored as (A, B, c2p) and as (B, A, p2c). The UCLA data set contains 368,048 links, i.e., 184,024 unique links of which each is stored in a forward and a backward representation. This is important while comparing the inferred data set and UCLA's data set. The advantage of this forward and backward representation is that similarities between the two data sets can be analyzed precisely.

The examined UCLA data set was created on 06/09/2013 covering the month of August. 128,462 (34,90%) of their observed AS links are labeled as unknown. Regarding c2p and p2p it shows almost the same ratio as our inferred data set, however, only half as many of the relationships were inferred by UCLA (see Figure 8).

To determine the exact differences between the two data sets, it is helpful to consider the entire set of links (forward and backward representation). Thus, every edge between two nodes A, B occurs twice, i.e., as A-B and B-A. Since links can be matched independently from their stored direction, differently inferred relationships can be compared precisely. The left pie chart in Figure 9 shows four categories. Category "same" stands for every link that occurs in both data sets and has the same relationship in both. The subset of links, which occur in both data sets but with different inferred relationships, is represented by "differs". Links that occur only in the inferred data set but do not appear in UCLA's data set belong to "additional". On the opposite, links that occur in UCLA's data set but do not appear in the

Figure 8: Number of unique links by relationship types of our inferred data set versus UCLA's data of August 2013. The inferred data set shows a ratio of 75% c2p relationships and 25% p2p relationships. UCLA's data set contains only half as many inferred business relationships. The ratio of 74% c2p to 26% p2p links, however, is similar to the inferred data set. Around one third of the UCLA links are labeled as unknown.

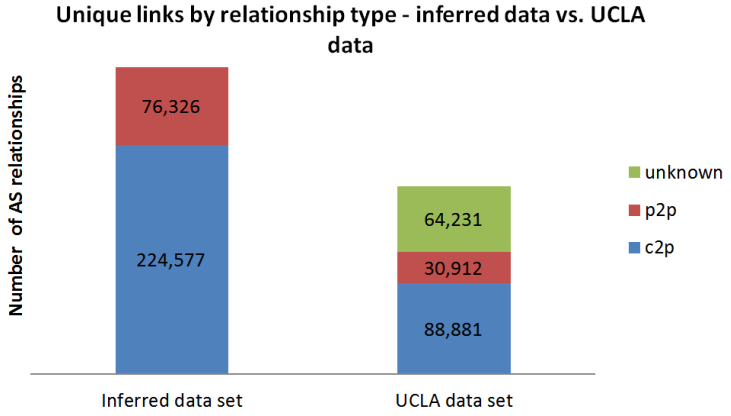

inferred data set are categorized "missing". In total 699,136 links or 349,568 unique links, respectively, can be observed by combining both data sets.

Figure 9 shows that 39% (i.e., differs+same) of all observed links occur in both data sets. Approximately two thirds of these carry the same business relationship, and for one third the two data sets show a different business relationship. The remaining 14% of all observed links are missing in the inferred data set. They account for 97,330 links. The pie on the right in Figure 9 gives an insight into how the missing links are labeled in the UCLA data set. It shows that the major part ($\approx$ 75%) of the missing links are labeled to be unknown. Broken down to unique links, only 3,458 customer-provider relationships and 10,230 peer-to-peer relationships inferred by UCLA are missing in our inferred data set. Together these account for only 3.9% of all observed unique links. Since most of the missing links are not labeled with any business relationship, these links do not contribute to the Internet graph of business relationships. The links of the small set of UCLA's business relationships that do not appear in our data set could be missing for several reasons, e.g., discarded paths during sanitizing or different BGP path sources that were used, and most likely the usage of other sources by UCLA.

Links which are labeled differently in the two data sets are illustrated in more detail in Figure 10. 95,810 ($\approx$ 14%) of all observed links carry different relationship labels in the two data sets. Again, a major part ($\approx$ 61%) of

Figure 9: Comparison of relationships by UCLA and inferred data set – focus on undiscovered relationships. In total 699,136 links or 349,568 unique links, respectively, were observed by combining both data sets. Approximately 75% of links which are missing in our inferred data set are labeled as unknown.

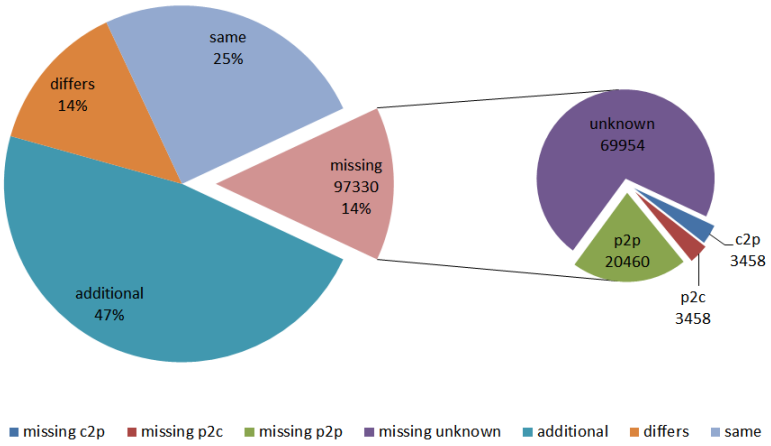

these links are labeled to be unknown in the UCLA data set. Considering unique links, 10,534 peer-to-peer relationships and 8,117 customer provider relationships are allocated by UCLA. This accounts for $\approx 5.3\%$ of all observed links.

The unknown links do not provide any hint on the correctness of the inferred data set. However, the small number of relationships that have been inferred differently supports the observation of Luckie et al. [14]. They observed by validation a performance of UCLA's algorithm close to their own algorithm. According to their investigation, UCLA's algorithm tends to produce errors by inferring c2p where the true business relationship is p2p. In such cases, the customer often has a larger degree than the provider. UCLA's approach assumes that over time every c2p relationship can be seen from a tier-1 AS. Any other observed link not visible from tier-1 is assigned p2p. However, reasons such as traffic engineering or selective announcements are causing a growing number of region-specific c2p relationships visible only below the provider AS. This leads to too many p2p inferences in the UCLA data set. Luckie et al. [14] take these effects for example in inference step 6 (see Section 4.2) of their algorithm into account.

In order to ensure the inferred data set has similar properties and therefore Luckie et al.'s algorithm has been implemented correctly, Figure 11 sheds light on the fraction 5.3% of differently inferred relationships. The pie chart shows that the main part ($\approx 57\%$) of these links are those assigned

Figure 10: Comparison of relationships by UCLA and inferred data set – focus on mistakenly inferred by UCLA. In total, 699,136 links or 349,568 unique links, respectively, were observed by combining both data sets. Approximately 61% of links which where inferred differently are labeled to be unknown by UCLA. Around 17% were inferred to have a customer-provider relationship, and the remaining 21% are labeled p2p by UCLA.

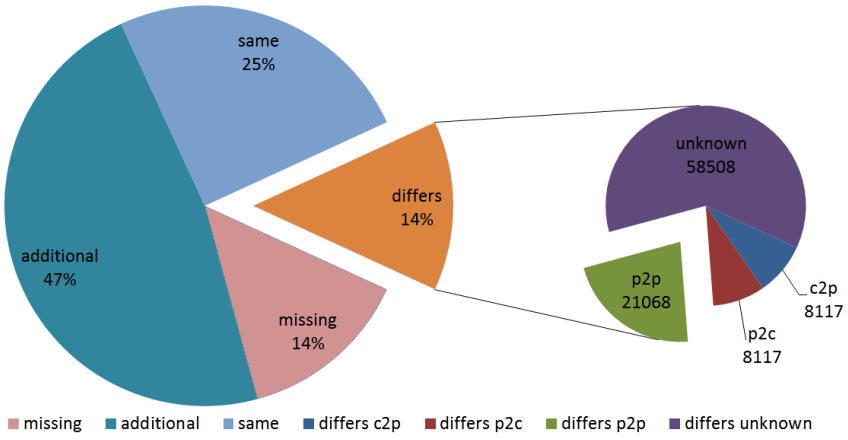

Figure 11: The figure shows the proportion of differently inferred business relationships between UCLA's and our inferred data set. The major part of these links are assigned c2p in the inferred data set and p2p by UCLA. The number of links with c2p (inferred data set) vs. p2c (UCLA) is negligible.

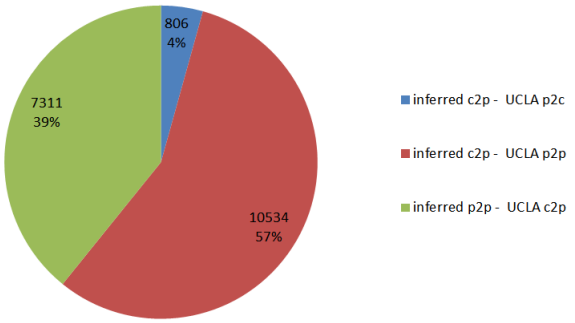

p2p by UCLA and c2p in the inferred data set. Furthermore, Luckie et al. mentioned that the number of occurrences where c2p was inferred by UCLA and the correct relationship was p2c is insignificant [14]. Only 806 of 349,568 unique relationships showed this combination.

5.2. AS Nodes

The presented implementation produced a total number of 56,048 unique ASNs. UCLA's data set of August 2013 contained 46,455 unique ASNs (see Table 1). In order to find reasons for links that were not seen in the inferred data set but in UCLA's data set, the assignment of discovered ASNs was checked. Therefore, the set of ASNs of the two data sets were compared with IANA's list of AS assignments (downloaded 16/09/2013) [32]. It turned out that the AS numbers of 439 ASes of UCLA's set were situated in number ranges that were not assigned to any RIR. Accordingly, paths containing these ASNs would have correctly been discarded and not considered for business relationship inference by the implemented algorithm. Looking at what these unassigned ASNs are reserved for (see Figure 12), it turns out that over 90% of these are reserved for private use.

Table 1: Number of assigned ASNs by Regional Internet Registry

Regional Internet Registry	**Inferred data set**	**UCLA**
Unassigned ASNs	0	439
Assigned by AFRINIC	810	694
Assigned by APNIC	7,007	5,700
Assigned by ARIN	20,367	16,336
Assigned by LACNIC	3,217	2,894
Assigned by RIPE NCC	24,647	20,391
Total number of ASes	**56,048**	**46,455**

Another intuition is that some paths would have been discarded if they contained ASNs which are actually IXPs since this condition was set in the implemented algorithm. Therefore, the assembled list of IXPs (see Section 4) was compared with all assigned ASes of UCLA's data set. Table 2 shows that 37 ASes are IXPs and would have been discarded by the implemented algorithm. Added to the number of unassigned ASNs, this sums up to a total number of 476 invalid ASNs in the UCLA data set. Accordingly, business relationships in UCLA's set where these ASNs participate should be invalid and could explain missing links (see 9) in the inferred data set. Thus, a total number of 2,322 UCLA relationships is invalid. Table 3 shows that of all UCLA links that are missing in the inferred data set 14% of c2p, 6% of p2c and 11% of p2p relationships can be explained by the usage of invalid

Figure 12: Unassigned ASNs within the UCLA data set. Over 90% of these are reserved for private use.

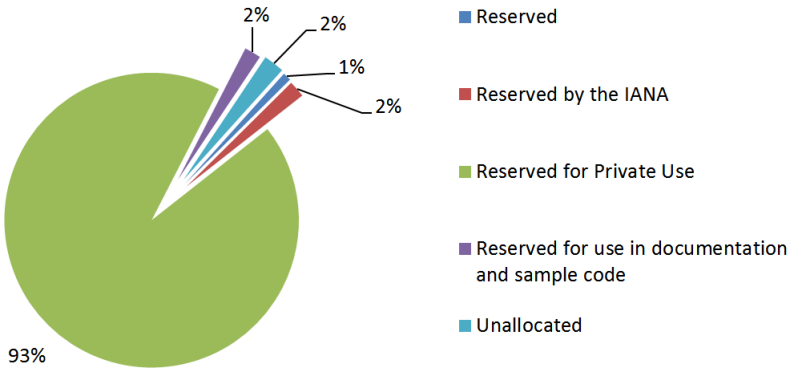

ASNs for inference. However, these are rather small proportions and do not explain the major part of missing links. Therefore, it can be concluded that the major part of missing links can be ascribed to the usage of additional BGP file sources by UCLA (e.g., Abilene [38]).

Table 2: Number of IXP and non-IXP ASes in UCLA data set

Regional Internet Registry	**Number of IXPs**	**No IXP**
AS_TRANS	1	0
Assigned by AFRINIC	2	692
Assigned by APNIC	4	5,696
Assigned by ARIN	4	16,332
Assigned by LACNIC	11	2,883
Assigned by RIPE NCC	15	20,376
Total number ASes	**37**	**45,979**

In order to determine the large number of additional links that occur in the inferred data set but not in UCLA's data set, a look at the VPs used could give some explanation. According to Oliveira et al. [9], UCLA used BGP data from ≈ 400 ASes provided by Route Views, RIPE-RIS [22], and Abilene [38]. In contrast, our inferred data set contains 595 VPs (see Table 4). Unfortunately, the exact ASNs of those ASes are not available. Hence, a direct mapping of VP ASes that do not occur in UCLA's data set to links that are categorized as additional is not possible.

However, the much higher number of used VPs clearly substantiates the observed difference in additionally inferred business relationships between the inferred data set and UCLA's data, at least to a certain extent.

Table 3: Number of missing relations with unassigned or IXP ASNs in UCLA data set

Description	**c2p**	**p2c**	**p2p**	**unknown**
AS_TRANS	96	14	25	151
Reserved	6	2	1	16
Reserved by the IANA	2	0	2	4
Reserved for Private Use	224	76	160	722
Reserved for use in documentation and sample code	3	0	4	8
Unallocated	4	0	1	9
IXP ASN	159	114	178	341
Sum of invalid relations	494	206	371	1251
Proportion of "missing" relations	**14%**	**6%**	**11%**	**36%**

Table 4: Number of vantage point and non-vantage-point ASes within in the inferred data set.

Vantage point	**Number of ASNs**
No	55,453
Yes	595
Sum	**56,048**

5.3. Validation

In the manuscript of Luckie et al. [14] it was somewhat unclear what time range for a single business relation graph inference had been used. Eventually, they provide monthly inference of a 5 day interval. Assembling validation data for business relationship inferences is like the inference itself not a trivial task. The validation data provided by Luckie et al. is based on corrected, previously inferred data sets from January 2010 and 2011 that had been inferred by using the algorithm of Dimitropoulos et al. [28]. As mentioned earlier, this algorithm inferred too many p2p relationships. Luckie et al. [14] corrected the data set by using three sources: directly reported data, RPSL data, and data from communities.

Directly reported data has been collected by feedback from network providers through CAIDA's website or via direct e-mail exchange. Furthermore, they collected data from routing policies which are stored by network operators in public databases. This data is stored in the Routing Policy Specification Language (RPSL). The routing policy of an AS is stored in the *aut-num* record which lists the export and import rules for each neighbor AS. Luckie et al. [14] used the RIPE WHOIS database, one of the largest sources for RPSL data, to derive their validation data.

The third source used by Luckie et al. in order to assemble the validation data set involves community attributes included in route announcements.

These are additional attributes that can be carried through several ASes propagating the route. However, the policy of an AS could also include removing this information from routed packets. These community attributes are often used by ASes to publicly document the meaning of their routing policies. Dictionaries with the meaning of individual community attributes are often provided on the websites of network operators and IRR databases.

Luckie et al. [14] mention that even these apparently reliable sources of business relationship information agree only on 99% of the links in the validation data set. To construct the validation data set, the sources were combined in the following precedence order: Directly reported data using CAIDA's website, RPSL data, BGP community data and directly reported information via e-mail exchange where e-mail exchange information has the highest priority. In conclusion, their validation data contains 41,604 relationships (16,248 p2p and 23,356 p2c) [14].

In order to prove correctness of our inferred data set and implementation, the inferred relationships have been directly compared to the relationships in the validation data set and grouped by relationship (see Table 5).

Table 5: Mapping of relationships between inferred data set and validation data set.

Inferred relationship	**True relationship**	**Number of relationships**
c2p		224,341
c2p	p2c	39
c2p	p2p	197
p2c	p2c	29,933
p2c	p2p	528
p2p	p2c	156
p2p	p2p	9,699

From Table 5, a confusion matrix has been derived for both types of relationships – p2p and p2c relationships (see Figure 13). These matrices build the basis to calculate PPV and TPR:

- $Recall = TPR = \frac{TruePositives}{ActualPositives} = \frac{TP}{TP+FN}$
- $Precision = PPV = \frac{TruePositives}{PredictedPositives} = \frac{TP}{TP+FP}$.

Table 6 shows TPR and PPV for p2p and p2c inferences of Luckie et al.'s validation and other algorithms they validated (i.e., [11, 28, 24]). It also includes TPR and PPV values of the inferred data set. The results show that both implementations performed close to each other. The most significant difference can be observed in the TPR of p2p inferences. The TPR of the

Figure 13: Confusion matrices for p2c and p2p inferences of the inferred data set and the validation data set provided by CAIDA.

p2c		Predicted Class	
		p2p	other
Actual Class	p2p	29933	195
	other	528	9699

p2p		Predicted Class	
		p2p	other
Actual Class	p2p	9699	725
	other	156	29972

inferred data set is 6.3% lower than Luckie et al.'s validation value. A brief examination of what causes the difference showed that the main cause were too many p2c relationships where the true relationship is p2p. This accounts for $\approx 75\%$ of the false negative values in the inferred data set.

Table 6: TPRs and PPVs of different relationship inference algorithms (data taken from [14]). TPR and PPV of the inferred data set shows similar performance as Luckie et al.'s own implementation.

Algorithm	c2p		p2p	
	PPV	**TPR**	**PPV**	**TPR**
Inferred data set	**98.4%**	**99.4%**	**98.4%**	**93.0%**
Luckie et al.	**99.6%**	**99.3%**	**98.7%**	**99.3%**
UCLA	99.0%	94.7%	91.7%	98.8%
Xia+Gao	91.3%	98.6%	96.6%	81.1%
Gao	82.9%	99.8%	99.5%	62.5%

Furthermore, Table 6 also shows the high level at which most algorithms perform. Improvements in TPR and PPV values can only be achieved within percentage fractions. Gao's algorithm [11] achieves the highest PPV for p2p relationships since it makes the fewest number of p2p inferences compared to other algorithms. However, that comes at the price of many more c2p inferences than actually exist [14].

In summary, this section indicates the correctness of the implementation of the used algorithm. A comparison has been drawn between the UCLA data set of business relationships and the inferred data set. Firstly, the results show that the inferred data set misses only a small fraction of the relevant links and their business relationships of UCLA's data. Furthermore, that subset could be partly explained by the fact that UCLA used ASes for their inference that would be considered invalid by the implemented algo-

rithm. However, the major part of missing relevant links can be ascribed to the use of additional BGP sources. Secondly, the analysis of differently inferred relationships shows that the major part of differently inferred relationships resolved to c2p when UCLA inferred p2p. This result conforms with the observations of Luckie et al. [14] and, thus, can be interpreted as evidence for a correct implementation.

6. Discussion

For the last decade researchers focused on the creation of Internet topologies at the AS level. Due to the common conception that participants have to pay each other for transit services different scientific peer groups tried to infer these business relationships and assemble a complete directed graph of the AS-level Internet. In the previous sections, the most promising algorithm in terms of correctness (by Luckie et al.) has been explicated. The result of this work is a directed graph representing the Internet on AS-level over the last 13 years. It has been shown that the properties of the inferred data set give strong evidence that the algorithm has been implemented correctly. Comparison with a validation data set assembled by Luckie et al. [14] has yielded correctness of the inferred data set and the implementation. Links belonging to categories "differs" and "missing" deliver most information for an interpretation of the inferred data set.

Category "missing". It has been shown that missing links (14%) that are actually allocated business relationships, account for only 3.9% of all observed links. The major part of missing links is labeled to be unknown by UCLA. Those missing links can partly be explained by unassigned ASes and IXP ASes which were not filtered by the UCLA algorithm (see Table 4). However, the majority of missing links can be ascribed to the use of additional BGP table sources. UCLA not only used Route Views as a source for MRT files but also RIPE-RIS and Abilene. In conclusion, by using additional source files the inferred topology of the Internet could be extended even further.

Category "differs". Among differently labeled links those which are unknown by UCLA ($\approx 61\%$) give no information for interpretation. However, those inferred differently confirm the observations Luckie et al. [14] made. Supposedly mistakenly inferred p2p by UCLA and inferred to be c2p in our data set account for 57% of differently inferred relationships. Obviously, the number of c2p links inferred by UCLA which are labeled p2p in the inferred

data set is with a ratio of 39% of differently inferred links also fairly high. In addition to the effects of traffic engineering and selective announcements, the much smaller time range and additional sources used by UCLA's explain this observation.

UCLA inferred MRT files from a time range of "only" $\approx$ 7 months. For our inferred data set, a time range of nearly 13 years of BGP routing files has been covered. Overall 972 RIB files of 13 different Route Views archives were collected. The higher number of used VPs ($\approx$ 400 UCLA vs. 595 our inferred data set) reflects the difference. This clearly comes with consequences while comparing the two data sets.

Validation. Comparing TPR and PPV for p2p and p2c inferences of Luckie et al. [14]'s validation results and the inferred data set has shown that, besides the TPR of p2p inferences, all values differ by less than 1%. The difference of $\approx$ 6% in TPR of p2p inferences were mainly caused by too many p2c inferences of relationships that are actually p2p relationships. One explanation for this result could be that the validation data reflect a recent state of the true business relationships on AS-level. According to Luckie et al. their validation data reach back to 2010. The inferred data set, however, used BGP data sources reaching back to 2001. Therefore, it is likely that some ASes occurred as a customer rather than a peer in an earlier state. Accordingly, they would appear in different positions in BGP RIB files in the past than nowadays, for instance if they had been announced as a customer in an older BGP file. However, the task to investigate this observation further will be left to future work. Despite the difference in TPR of p2p inferences, the close performance of the presented implementation and Luckie et al.'s data confirm correct inference and implementation.

Projects such as UCLA and CAIDA have so far only provided snapshots that only capture short time ranges reaching from a few days to a few months. The presented inference certainly provides a better perspective of the entirety of the Internet topology by comprising a view over the past 13 years. For the implementation of Luckie et al.'s algorithm the set of IXP has been extended by using data provided by Augustin [33]. Furthermore, ASes in ASDOT format have been converted to ASPLAIN. Paths containing ASes in ASDOT format could, therefore, been used for relationship inference. Another advantageous aspect of the presented implementation is the usage of a provider cone instead of a customer cone for this implementation. It caused linear runtime and moderate memory usage during the inference process.

6.1. Limitations

While a long time range gives a comprehensive view of the AS-level topology, the drawbacks of this setting should be considered. In the following subsection the major shortcomings will be presented.

Reassignment of ASNs. One question that came up during examination of the inferred data set is: Are ASNs generally reused and does this have an influence on inferring business relationships? The intuition behind this question is the following: If an ASN X is assigned to a certain AS A at an earlier point of time and to another AS B later, the two distinct ASes would be considered as one domain. Each AS would have certain links and business relationships to other ASes and by that a certain position in the Internet topology. Accordingly, since those ASes are identified by their ASN they would be considered as one AS. As a consequence, the two sets of links to X would be merged. Since the two ASes could appear on different hierarchy levels, the inferred relationships involving this ASN would not reflect a picture of reality.

In order to answer this question, the assignment and advertisement of ASNs has to be illuminated. On top of this system stands the Internet Assigned Numbers Authority (IANA) which administrates the pool of ASNs. IANA does not assign ASNs directly to ISPs but allocates ASNs to Regional Internet Registries (RIRs). The RIRs, on the other hand, assign ASNs to different ISPs which, in turn, advertise the ASNs in routing tables [32].

In order to generate projections of ASN consumption, Huston [8] explored ASNs in more detail. He stated that under current conditions, costs for maintaining an ASN allocation do not exist. Once an allocation of ASNs by IANA to an RIR is made, it remains in this state. There is no incentive for RIRs to return ASNs [8] Huston [8] observed that the growth rates in the number of unadvertised ASes (i.e., never announced in a BGP file) is slightly lower than the growth rate of the number of advertised ASes over time. The ASN pool is being consumed in a numerical sequence, i.e., more recently allocated ASNs are higher than older ones. Looking at this distribution of ASNs in blocks of 256, it turns out that older blocks (i.e., blocks with lower ASNs) contain higher proportions of unadvertised AS than more recent blocks. Figure 14 shows the breakdown of each block of 256 AS numbers, looking at how many numbers from each of these smaller pools is currently advertised, how many are not advertised, and how many are in the RIR pre-allocation pools [8].

The age distribution of unadvertised ASNs (see Figure 15) helps to understand the meaning of this observation. It appears that the probability of

Figure 14: ASN Status by Block. The figure shows the breakdown of each block of 256 AS numbers, looking at how many numbers from each of these smaller pools is currently advertised, how much are not advertised, and how much are in RIR pre-allocation pools (cited from [8]).

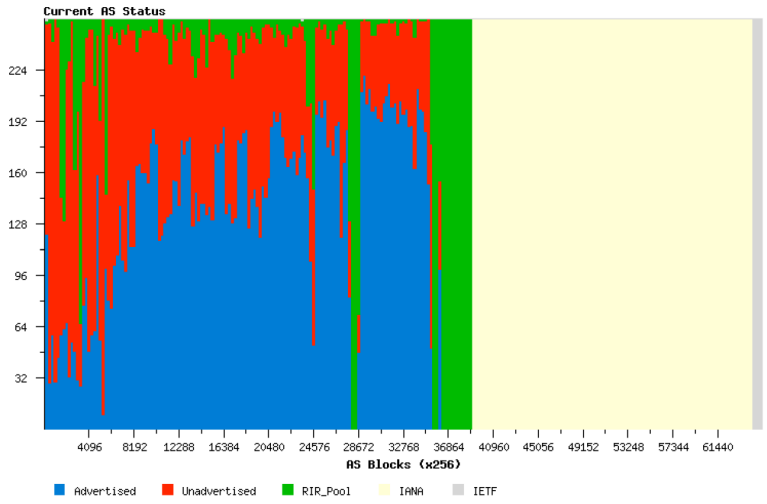

an ASN being visible in the routing table is directly associated to the time that has passed since allocation. Furthermore, the figure shows a peak in the number of unadvertised ASNs that have the latest allocation date. It shows that there's a three-month gap between the allocation of an ASN and its advertisement in the BGP routing tables. Considering all the foregoing observations, it can be concluded that an ASN which is no longer needed does not get returned to the pool of unallocated ASNs for later reuse. Instead, it is set to a latent state where the AS is still considered allocated but unadvertised [8]. This leads to the conclusion that in the chain of *IANA – RIR – ISP* ASNs are generally not recycled and reused.

Stable versus unstable paths?. The presented algorithm explicitly uses unstable and stable paths for inference. Dimitropoulos et al. [28], in contrast, used heuristics which evaluate whether or not a path occurs in all 15 sets of BGP tables which are divided by 8-hour intervals. A path is considered unstable if it does not occur in each of the sets and stable if it does. Unstable paths are considered to be caused by misconfigurations [28]. However, Luckie et al. [14] argue that backup links are more likely to be included if all AS paths are used, and discarding of normally stable paths due to temporary peering disputes is prevented [14].

On the other hand, Mahajan et al. [39] found that misconfiguration

Figure 15: Proportion of unadvertised ASes by RIR Allocation Date [8].

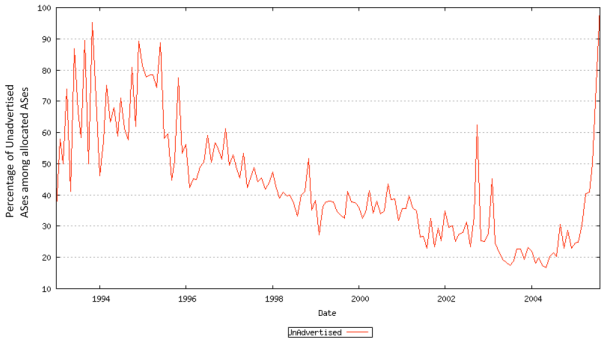

of routers is an issue. 200-1,200 BGP table entries are affected each day [39]. Dimitropoulos et al. [28] developed the algorithm which was formerly used by CAIDA for business relationship inference. In their paper they presented a heuristic to account for misconfiguration in router paths during the sanitizing phase. They collected routing data from 15 different BGP table instances. By deriving the persistence function from these 15 sets, it turned out that the majority of paths appears in most of the 15 sets. However, a significant number of paths appeared in only a subset of the 15 sets. Consequently, they used only stable paths for inference and observed only little loss of information. The set of stable paths was just 12.49% smaller. The observed reduction in number of links (4.34% less) and ASes (1.87% less) was rather insignificant.

Luckie et al. [14] address misconfiguration by a heuristic used in their algorithm. They minimize false c2p inferences by looking at the distance of an AS to a VP. This heuristic accounts for the problem of ASes leaking provider or peer routes to peers. It is based on the intuition that, in order to become globally reachable, each AS enters a provider relationship. Accordingly, at least one VP should see the provider announcing the customer's routes.

6.2. Outlook on Further Research

Relationship inference for the AS-level Internet has reached its limits of accuracy due to more and more sophisticated approaches and the fact that there is no 100% accurate ground-truth of the AS-level topology. However, it has been shown that it is possible to achieve a more comprehensive AS-level map by using extensive BGP sources. The implemented approach focused on simple business relationships (p2c and p2p). However, actual relationships

between ASes are often more complex. Further research could be conducted to infer relationships such as sibling-to-sibling or other hybrid relationships.

Another starting point for future research could be the measure for ranking AS nodes. Luckie et al. [14] already used a measure called transit degree instead of the rather simple node degree. However, there exists a fair amount of more sophisticated topological metrics that could potentially be used to rank AS nodes [2][40][4][5].

7. Conclusion

The Internet, as a global critical infrastructure, needs to be studied with respect to robustness and resilience as well as power structure and information flows. All of these important investigations need a solid understanding and mapping of the Internet topology as a vast and complex routing topology but also business infrastructure. Since the Internet has a decentralized structure, there are no accurate and comprehensive maps of the Internet readily accessible. This article presents one implementation of the most recent and most promising approach for relationship inference on the AS level. The algorithm has been improved in terms performance and quality of the sanitizing process. Unlike recent projects, not a only snapshot of the topology of the Internet has been inferred but a comprehensive map showing the Internet over the last decade. The correctness of this implementation and the inferred data set was examined by comparison with a business relationship graph and a validation data set.

8. Acknowledgments

The authors would like to thank Annika Baumann for very valuable discussions.

References

[1] CNNMoney, . Fortune 500. 2013. URL `http://money.cnn.com/magazines/fortune/fortune500/2013/`.

[2] Baumann, A., Fabian, B.. How Robust is the Internet? – Insights from Graph Analysis. In: Proceedings of the 9th International Conference on Risks and Security of Internet and Systems (CRiSIS 2014), Trento, Italy. LNCS 8924; Springer; 2014,.

[3] Baumann, A., Fabian, B.. Towards Measuring the Geographic and Political Resilience of the Internet. International Journal of Networking and Virtual Organisations 2013;13(4):365–384.

[4] Baumann, A., Fabian, B.. Vulnerability Against Internet Disruptions – A Graph-based Perspective. In: Proceedings of the 10th International Conference on Critical Information Infrastructures Security (CRITIS 2015), Berlin, Germany. LNCS; Springer; 2015,.

[5] Fabian, B., Baumann, A., Lackner, J.. Topological Analysis of Cloud Service Connectivity. Computers & Industrial Engineering 2015;88:151–165.

[6] Baumann, A., Fabian, B.. Who Runs the Internet? Classifying Autonomous Systems into Industries. In: Proceedings of the 10th International Conference on Web Information Systems and Technologies (WEBIST), Barcelona, Spain. 2014,.

[7] Haddadi, H., Rio, M., Iannaccone, G., Moore, A., Mortier, R.. Network topologies: Inference, modeling, and generation. IEEE Communications Surveys & Tutorials 2008;10(2):48–69.

[8] Huston, G.. Exploring autonomous system numbers. The ISP Column 2005;URL `http://203.133.248.2/ispcol/2005-08/as.pdf`.

[9] Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.. The (In)Completeness of the observed internet AS-level structure. IEEE/ACM Transactions on Networking 2010;18(1):109–122.

[10] Roughan, M., Willinger, W., Maennel, O., Perouli, D., Bush, R.. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. IEEE Journal on Selected Areas in Communications 2011;29(9):1810–1821.

[11] Gao, L.. On Inferring Autonomous System Relationships in the Internet. IEEE/ACM Transactions on Networking (ToN) 2001;9(6):733–745.

[12] CAIDA, . AS relationships. 2013. URL `http://www.caida.org/data/active/as-relationships/index.xml`.

[13] UCLA, . IRL topology. 2013. URL `http://irl.cs.ucla.edu/topology/#data`.

[14] Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., kc claffy, . AS Relationships, Customer Cones, and Validation. In: Proceedings of the 2013 Conference on Internet Measurement. ACM; 2013, p. 243–256.

[15] Christiansen, T., Wall, L., Foy, B., Orwant, J.. Programming Perl: Unmatched Power for Text Processing and Scripting. O'Reilly Series; O'Reilly Media, Incorporated; 2012. ISBN 9780596004927.

[16] Morin, R., Brown, V.. Scripting languages. 2013. URL http://www.mactech.com/articles/mactech/Vol.15/15.09/ScriptingLanguages/index.html.

[17] Route Views, . Route views tools page. 2013. URL http://www.routeviews.org/tools.html.

[18] Witten, I., Frank, E., Hall, M.. Data Mining: Practical Machine Learning Tools and Techniques: Practical Machine Learning Tools and Techniques. Elsevier Science; 2011.

[19] Rekhter, Y., Li, T., Hares, S.. RFC 4271: A border gateway protocol 4 (BGP-4). 2006. URL http://tools.ietf.org/html/rfc4271.

[20] Gundalwar, P.R., Chavan, V.N.. Border Gateway Protocol (BGP): A Simulation Based Overview. International Journal of Application or Innovation in Engineering & Management (IJAIEM) 1999;.

[21] Gregori, E., Improta, A., Lenzini, L., Rossi, L., Sani, L.. On the Incompleteness of the AS-level Graph: A Novel Methodology for BGP Route Collector Placement. In: Proceedings of the 2012 ACM Conference on Internet Measurement. 2012, p. 253–264.

[22] RIPE-RIS, . Routing information service (RIS) — RIPE network coordination centre. 2013. URL http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[23] Blunk, L., Karir, M., Labovitz, C.. RFC 6396: Multi-threaded routing toolkit (MRT) routing information export format. 2013. URL http://tools.ietf.org/html/rfc6396.

[24] Xia, J., Gao, L.. On the Evaluation of AS Relationship Inferences. In: Proceedings of IEEE GLOBECOM'04; vol. 3. 2004, p. 1373–1377.

[25] Subramanian, L., Agarwal, S., Rexford, J., Katz, R.H.. Characterizing the Internet Hierarchy from Multiple Vantage Points. In:

Proceedings of the 21st IEEE INFOCOM Conference; vol. 2. 2002, p. 618–627.

[26] Aldous, J., Wilson, R.. Graphs and Applications: An Introductory Approach; vol. 1 of *Graphs and Applications: An Introductory Approach*. Dickenson; 2000. ISBN 9781852332594.

[27] Di Battista, G., Patrignani, M., Pizzonia, M.. Computing the Types of the Relationships between Autonomous Systems. In: Proceedings of the 22nd IEEE INFOCOM Conference; vol. 1. 2003, p. 156–165.

[28] Dimitropoulos, X., Krioukov, D., Fomenkov, M., Huffaker, B., Hyun, Y., Riley, G.. AS relationships: Inference and validation. ACM SIGCOMM Computer Communication Review 2007;37(1):29–40.

[29] Wang, H., Xie, H., Yang, Y.R., Li, L.E., Liu, Y., Silberschatz, A.. Stable Egress Route Selection for Interdomain Traffic Engineering: Model and Analysis. In: Proceedings of the 13th IEEE International Conference on Network Protocols (ICNP 2005). 2005,.

[30] Zhang, B., Liu, R., Massey, D., Zhang, L.. Collecting the Internet AS-level Topology. ACM SIGCOMM Computer Communication Review 2005;35(1):53–61.

[31] Valiente, G.. Algorithms on Trees and Graphs. Springer; 2002.

[32] IANA, . Autonomous system (AS) numbers. `http://www.iana.org/assignments/as-numbers/as-numbers.xml`; 2013.

[33] Augustin, B.. IXP mapping project. 2012. URL `http://www-rp.lip6.fr/~augustin/ixp/`.

[34] Potaroo, . AS names. 2013. URL `http://bgp.potaroo.net/cidr/autnums.html`.

[35] Bron, C., Kerbosch, J.. Algorithm 457: Finding all cliques of an undirected graph. Communications of the ACM 1973;16(9):575–577.

[36] Zhang, Y., Oliveira, R., Zhang, H., Zhang, L.. Quantifying the Pitfalls of Traceroute in AS Connectivity Inference. In: Passive and Active Measurement. 2010, p. 91–100.

[37] Augustin, B., Krishnamurthy, B., Willinger, W.. IXPs: Mapped? In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement. 2009, p. 336–349.

[38] Abilene, . The abilene observatory data collections. 2013. URL http://www.internet2.edu/network/observatory/data-collections.html/.

[39] Mahajan, R., Wetherall, D., Anderson, T.. Understanding BGP Misconfiguration. In: ACM SIGCOMM Computer Communication Review; vol. 32. 2002, p. 3–16.

[40] Mahadevan, P., Krioukov, D., Fomenkov, M., Dimitropoulos, X., kc claffy, , Vahdat, A.. The Internet AS-level Topology: Three Data Sources and One Definitive Metric. ACM SIGCOMM Computer Communication Review 2006;36(1):17–26.