Technische Universität Braunschweig

IAS | INSTITUTE FOR APPLICATION SECURITY

Bachelor's Thesis

Mapping and Analyzing BGP-level Hubs of Control in the Global Internet

Tim Christian Sauer

January 25, 2022

Institute for Application Security
Prof. Dr. Martin Johns

Supervisor:
Simon Koch, M. Sc.

Statement of Originality

This thesis has been performed independently with the support of my supervisor. To the best of the author's knowledge, this thesis contains no material previously published or written by another person except where due reference is made in the text.

Braunschweig, January 25, 2022

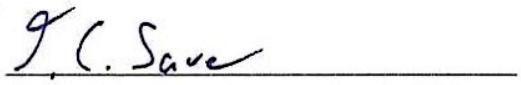

Abstract

The Internet is a network of networks. Those *subnetworks* are called *Autonomous Systems (ASes)*. For pathfinding of *AS-paths* between ASes the *Border Gateway Protocol (BGP)* is used. This thesis provides an introduction to BGP, its vulnerabilities and existing research. Furthermore we present a methodology to analyze which ASes are heavily traversed by AS-paths. Those ASes are considered *hubs of control*. We investigate into these hubs and show that a small fraction of them is present on nearly all of the AS-paths among the Internet. This fact raises attention, since existing research shows that attackers that have access on a high amount of paths are even able to unravel users of the privacy network *Tor*. Our analysis shows, which ASes act as hubs for *local topology*, e.g. for certain countries and which act as hubs for *global topology* and therefore for the entire Internet. We are the first ones who run an analysis not only for todays topology, but also take historical data into account and infer trends over the years. On top of that, we investigate, if there are ASes from the evaluated hubs that cannot be circumvented in routing and are therefore essential.
One of our key findings is, that there is indeed a small fraction of ASes which is represented on the worlds AS-paths. We also identify the tendency, that the total amount of ASes around the world is increasing while the number of hubs of control stays the same. Therefore few hubs of control will increase their power over the next years. But not all hope is lost: In view of the fact, that there are few essential ASes, the stability of the Internet is not dependent on single systems.

Contents

List of Figures

List of Tables

1. Introduction

The Internet is a network of networks. Those subnetworks are called *Autonomous Systems (ASes)*. Between Autonomous Systems the *Border Gateway Protocol (BGP)* is used for pathfinding. As BGP was first designed in 1990 there have been little efforts towards security at this time. Over the years there have been more and more incidents of hijack and even interception of traffic, based on exploiting flaws in BGP.
This is especially alarming, as we live in a time of permanently increasing globalization. In this process the Internet is *the* essential technology that connects people all over the world. One consequence of this ongoing process is, that data becomes one of the most valuable goods of our time. Several institutions all over the world permanently collect, categorize and analyze data. Given this fact, it becomes a question where data goes when it leaves homes, organizations and countries. As the Internet does not have any physical borders, data often traverses multiple countries on the way to its destinations. It even appears that traffic that has its origin and destination in the same country leaves it on its way [1, 2]. This brings up the question, which countries data travels and if those countries are assumed trustworthy by the data's sender. The more halts data takes on its path, the more opportunities exist for a potential eavesdropper. Existing research underlines the severity of this topic by showing: An eavesdropper who captures a high enough amount of traffic is even able to unravel users of the well known *Tor network* [3] which aims to provide anonymity for its users.
Let us assume a country whose government aims to spy on its citizens traffic. To do so it would have to attack those ASes which are used for routing the citizens traffic.
One could question if such attacks are even necessary if those targeted ASes are already under the control of the given countries government. This brings up our research question: Let there be a set A_{orig} of origin ASes and a set B_{dest} of destination ASes. AB is the set of all existing paths between all origins from A_{orig} and all destinations from B_{dest}.

Are there certain ASes which are along a high number of AS-paths from AB and therefore state interesting points for observation of Internet traffic?

To approach the research question this thesis will give the following contributions:

1. We document the state of knowledge regarding the analysis of AS connectivity and important local and global hubs.

2. We present a methodology to analyze ASes from given input sets A_{orig} and B_{dest} and the resulting AS-paths from AB. ASes which are intensively used along AS-paths from AB and therefore state potentially sensitive points for observation are subjects of further analysis.

3. We geolocate ASes to monitor which ones of them are important hubs for local (inside a country) and which ones for global connectivity (between different countries).
4. We investigate if there are *essential* ASes, from those hubs discovered in point Contribution 3, which cannot be circumvented when routing traffic.

The Internet undergoes constant and rapid growth and change, potentially rendering observational work obsolete in few years. Contribution 2 and Contribution 3 revisit existing work [4] and analyze the changes since the last time, measurements were taken. As the number of ASes since then have doubled [5], significant changes in topology can be expected.

This thesis is structured as follows:

- We will cover basic concepts of BGP and introduce several attack types that exploit BGP's architecture.
- Subsequently, we will give an overview of previous work, that is related to this thesis.
- We introduce our methodology of analyzing ASes and AS-paths to identify hubs and essential ASes.
- Furthermore, we describe, how we implemented our methodology and a tool to visualize analyzed ASes and paths.
- We will outline challenges and limitations of our work.
- Finally, we present the results of our analysis, discuss them and propose consequent steps that could be part of work, done in the future.

2. Background

This section covers basic principals of routing in the Internet, implemented by BGP. Furthermore we give insights into past events where AS-level configurations went wrong and lead to loss of connections for a major fraction of Internet users. We will also cover the most common attack types, exploiting BGPs architecture.

2.1. Autonomous Systems and the Border Gateway Protocol

The Internet is a decentralized system built of a rapidly increasing number of routers. Several routers that share the same IP-prefix[1] and are under the administration of a common organization (e.g. Internet Service Provider (ISP), Scientific Institution) form an *Autonomous System* (AS, Plural: ASes). ASes are identified by a globally unique numerical identifier called *Autonomous System Number* (ASN).
Every device that is connected to the Internet is also part of an Autonomous System. Consider a person using its phone to order a product from *Amazon.com*. This user might have a phone contract from *Deutsche Telekom*. Consequently, its phone communicates via a *Telekom* router. The data it sends to order a product is routed from *Deutsche Telekom*s AS to *Amazon*s AS.
Inside of ASes, routers communicate based on *Interior Gateway Protocols*[2]. To the outside world they act as one system. Not every AS is physically connected to every other AS around the globe. Therefore ASes use paths, traversing 0 or more other ASes to communicate with each other. Those paths are called *AS-paths* or simply *routes*. To find AS-paths between ASes, the *Border Gateway Protocol* (BGP) is used. Section 2.4 gives details on routing between ASes and its implementation by BGP.

2.2. Matching IP-addresses to subnetworks

Every device that is connected to the Internet has an unique *IP-address* it can be identified by. Several sequential IP-addresses can be summarized to an *IP-subnetwork* (short *subnetwork*). Subnetworks can be identified by the IP-prefix of their IP-addresses. To find out which part of an IP-address is the prefix, the *subnetmask* is used. The subnetmask marks this part of an IP-address which is the IP-prefix. This is why an IP-subnetwork is always

[1] An IP-prefix is the part of an IP-address which identifies the network, a target device is located in. In other sources it might also be called *host-portion*.

[2] Interior Gateway Protocols (IGP) are used for routing inside of an Autonomous System. There exists a wide range of IGPs as every network is different and therefore has different requirements for a protocol.

given as a pair of an IP-address and a subnetmask.
Recall that every AS has one or more prefixes assigned. When traffic is send towards a certain destination, this destination is identified by its IP-address. To understand how an AS 'knows' which prefix and therefore which AS to send data to, we have to understand how to find out, which prefix an IP-address is part of. Consider Figure 2.1 for an example.

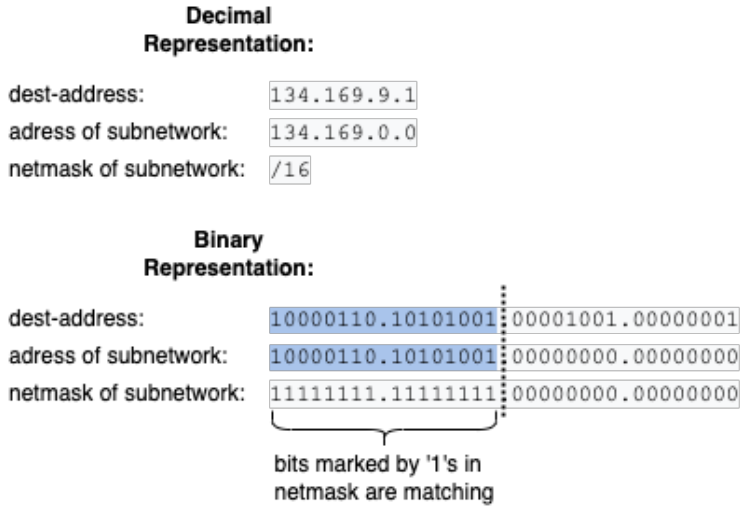

Figure 2.1.: Matching IP-addresses to their respecting IP-subnetwork.

The decimal representation shows the IP-address `134.169.9.1` and the address of an IP-subnetwork, `134.169.0.0`. Additionally, the IP-subnetwork has the subnetmask `/16`. To find out, whether the IP-address is part of the IP-subnetwork, the binary representations of the IP-address and the address of the IP-subnetwork have to be evaluated and are noted in octets. (See 'Binary Representation') Additionally, there are as many '1's noted, as the subnetmask is long. To fill the rest of the four octets, '0's are noted. It is now checked, whether the parts of the IP-address and the subnetwork-address which are marked by '1's are equal. If they are, the IP-address is located in the subnetwork, otherwise not.

2.3. Business Relationships between ASes

In order to understand, why Autonomous Systems take exactly the AS-paths they do, we have to know, that not all ASes act on a par with each other. The Internet is hierarchical structured and therefore ASes are assigned different tiers. In their paper from 2001 Gao et al. present a model of how ASes exchange data based on business relationships [6]. These relationships are caused by contracts between AS owners, which define pricing for traffic exchange. There are three roles ASes can take in relationship to other ASes:

1. Customer:
   A customer is usually a smaller AS that pays a larger AS for routing its traffic from or to other ASes.

2. Provider:
   A provider is a larger AS which is paid for routing a smaller ASes (customers) traffic.

3. Peer:
   Two ASes are peers to each other (peering ASes) if they have typically the same size and both benefit from routing traffic over each other and to each others clients.

Depending on their relationships to other ASes they can be classified into 3 tiers. ASes that only have customers or peers as their neighbors and can therefore reach the entire Internet free of charges are called tier-1 ASes. ASes that need both, peers and providers to reach the Internet are called tier-2 ASes. The third category are tier-3 ASes. Those systems only communicate via providers and therefore pay for all their routed traffic. Typical tier-1 ASes are ASes of ISPs which provide Internet service worldwide, without being charged. Some of them are *Deutsche Telekom*, *AT&T* and *China Telecom*. However most ISPs ASes are not able to reach the entire Internet for free and are therefore tier-2 ASes. Note, that not only ISPs can be tier-2 ASes. There are several ASes from sectors like Education/Research, Non-Profit, Government, etc. which purchase connectivity but also provide Internet access to other ASes and are therefore tier-2 ASes. For instance, AS680 which is the German research network (*DFN*) provides Internet access for German universities and other research institutions. At the same time, their Internet access is provided by *Deutsche Telekom*. Therefore AS680 is a tier-2 AS.

2.4. BGP-based routing

Fundamentally, BGP is a *Distant Vector Protocol*[3]. In contrast to other Distant Vector Protocols, BGP has to implement policies between different AS-owners. These policies can be conditions like: Traffic that has its origin in Germany should never be routed via the United States unless they are its destination.
Figure 2.2 shows how an AS A establishes a connection to a physically connected neighbor B, announces *UPDATE* messages to this neighbor and receives UPDATEs from it. This communication is done by routers at the edge of ASes, called *Border-routers*. ASes which have multiple Border-routers and therefore multiple connections to the Internet Backbone are called *multi homed ASes*. The following types of messages, called *BGP-announcements* can be send by the two communication parties:

message	**usage**
OPEN	Used to establish a connection. Contains protocol version, ASN, hold-timer...
UPDATE	Used to communicate new AS-paths or changes in existing AS-paths.
NOTIFICATION	Used to inform the other party about routing status, errors. Also used to close an established connection.
KEEPALIVE	Must be send by both parties before the hold-timer exceeds to hold connection for another hold-timer period.

The first step of routing is to establish a connection. Border-routers do not have to configure, which neighbors can notify them. The only way to prohibit communication to a specific neighbor would be to decline communication or cut the physical connection. To connect, AS As Border-router sends an *OPEN* message to its neighbors B Border-router. Bs Border-router confirms the connection with a *KEEPALIVE* message. After the establishment the connection is open for the period of one hold-timer. The two routers can send each other UPDATE messages. If both routers send a KEEPALIVE message before the hold-timer exceeds, the connection remains established, for the period of another hold-timer, otherwise it is closed. The two routers are also able to use *NOTIFICATION* messages to inform each other about routing status, errors, etc. A NOTIFICATION is also used to close a connection.
To propagate AS-paths, Border-routers use the UPDATE message. One UPDATE can hold several AS-paths. An AS-path is formed when an AS A sends its ASN and IP-prefix to a connected neighbor B via an UPDATE message. This AS-path now tells B via which hops

[3]Distance Vector Protocols use distance vector routing. Referring to Tanenbaum, distance vector routing is described as follows: 'In distance vector routing, each router maintains a routing table indexed by, and containing one entry for, each router in the subnet. This entry contains two parts: the preferred outgoing line to use for that destination, and an estimate of the time or distance to that destination.' [7]

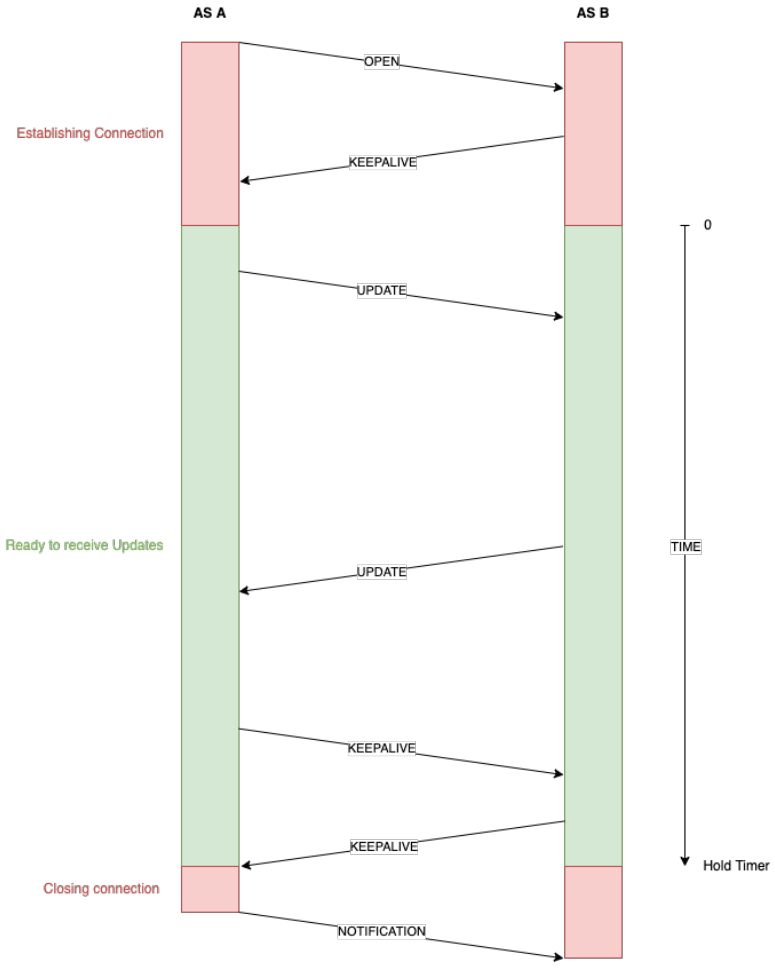

Figure 2.2.: Establishment and communication during a BGP-connection.

it can reach A. After an AS-path is first created by one AS, it only holds one hop, but as the path is later passed on by other ASes it grows in its length. B receives the AS-paths from As UPDATE and stores it to a database called *Routing Information Base (RIB)* or *Routing Table*. It happens quiet often, that B learns different AS-paths from its neighbors that lead to the same IP-prefix, at the end of the path. In this case it has to decide which path to actively use for routing. B now marks the preferred path as 'active' in its RIB. How tie-breaking between different paths works in detail is described in Section 2.4.1.
After choosing the preferred paths from its RIB and marking them as 'active' B prepends[4] its own IP-prefix and ASN to the AS-paths and propagates them on to its own neighbors via UPDATE messages. By doing the process repeatedly, an AS-path grows in its number of hops. Note, that an AS can choose, if it actually passes on received AS-paths. If it does, it is likely to happen that it has to route traffic for other ASes. As not every AS can route traffic free of charge via its neighbors, not every AS will pass on received AS-paths.

[4] **Prepending** ASN and prefix is the reason why the destination of an AS-path is always the last ASN (at the end of a path).

Figure 2.3 shows a part of a RIB, taken from a router at the *DE-CIX*[5]. Let this router be *R*1. Paths have been filtered for the AS of TU Braunschweig as origin (AS680). Information contained in the RIB are represented by columns. Every AS-path contained in the RIB is represented by one line of the table. As one can see, for all three entries in the RIB the 'Path' field only contains AS 680. Hence, *R*1 can reach TU Braunschweig without a hop in between. As the 'Status' field shows, the first route is the one that is currently marked as 'active' by this router. The 'Network' field shows the origin IP-subnetwork at the end of the path. Since the IP-address of TU-Braunschweig (`134.169.9.1`) is part of this IP-subnetwork (`134.169.0.0/16`) (recall Section 2.2 on how to find out), this path is a valid one for TU-Braunschweig as destination of traffic. The 'Next Hop' field represents the router to which *R*1 would send traffic routed towards TU-Braunschweig. We now combine information from the 'Learned' field with the 'Origin' field which says 'IGP' (which stands for Interior Gateway Protocol). Thus, we learn for all three entries, that the routes must stem from different sources from the inside of *R*1s AS (AS6939). 'LocPrf' stands for *Local Preference* which is one key attribute used for tie-breaking between different AS-paths with the same destination. The rest of the fields represent Cisco-internal parameters which are contained in the RIB as *R*1 uses Cisco-software.

core1.fra1.he.net> show ip bgp routes detail 134.169.9.1

Matching Routes	3
Status Codes	**A** - Aggregate **B** - Best **b** - Not Install Best **C** - Confederation eBGP **D** - Damped **E** - eBGP **H** - History **I** - iBGP **L** - Local **M** - Multipath **m** - Not Installed Multipath **S** - Suppressed **F** - Filtered **s** - Stale **x** - Best-External

Status	Network	Next Hop	Learned	Metric	LocPrf	Weight	Path	Origin	ROA
BI	134.169.0.0/16	80.81.192.222	216.218.253.18 (6939)	20	100	0	680	IGP	✅
I	134.169.0.0/16	194.146.118.60	216.218.252.79 (6939)	50	100	0	680	IGP	✅
I	134.169.0.0/16	193.178.185.42	216.218.252.78 (6939)	162	100	0	680	IGP	✅

Last Update 5d8h41m59s ago (1 path installed)

Entry cached for another 60 seconds. 2022-01-07 13:19:14 UTC

Figure 2.3.: Routing table dump from a router at DE-CIX. Paths are filtered for the AS of TU Braunschweig as destination.

After AS-paths are stored to the RIB and it has been decided which AS-paths are active, we now take a look at how routing of traffic is done. Consider Figure 2.4. Router *R*1 a Border-router of AS3801 gets an IP-package from the inside of its AS. It now checks whether it has a route in its RIB that is originated by the destination IP-address or an IP-subnetwork that includes it. It turns out that `134.169.9.1` is contained in the IP-subnetwork `134.169.0.0/16`. Therefore *R*1 passes the IP-package on to *R*2, a Border-router of AS174, as *R*2s IP-address is included in the 'Next-hop'-field of *R*1s chosen route. *R*2 receives the package from *R*1 and looks up the packages destination-address in its routing table. It finds a route which proposes `134.169.0.1` at AS680 as the next hop. Therefore it routes the package towards *R*3 to which this address belongs.
*R*3 receives the package and notices that the destination-IP-address is located inside its AS. Therefore it routes the package towards its destination via AS-internal routing.

[5]The DE-CIX is the largest Internet exchange point concerning the average throughput of traffic, worldwide. State January, 2022 it routes 6.9 Tbit/s [8].

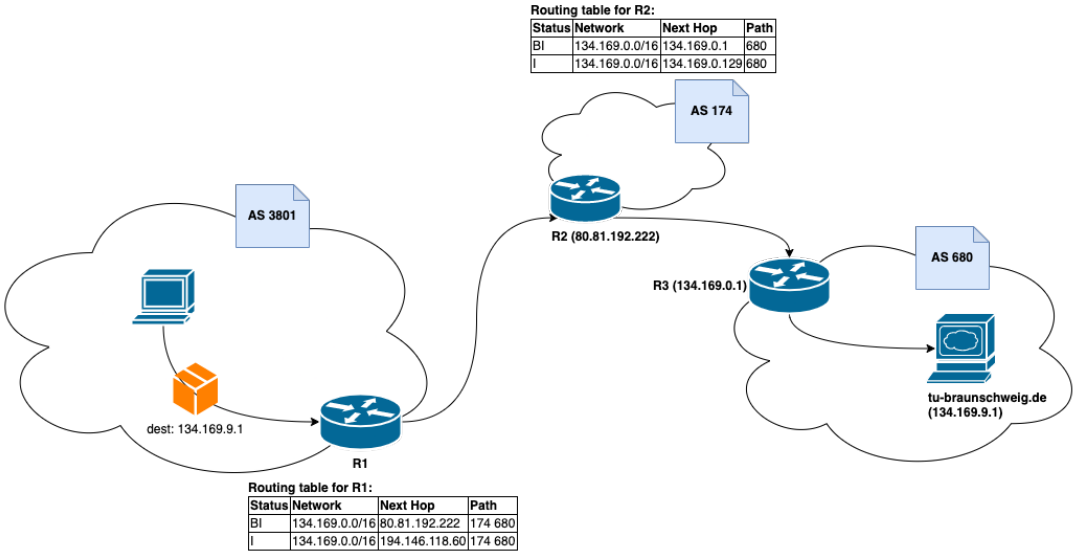

Figure 2.4.: Exemplary routing of an IP-package towards a destination.

Note that traffic always takes the opposite direction of AS-paths. If we talk about origins and destinations in the following chapters we always mean origin and destination in terms of AS-paths. Hence, if a path originating in A and destining B is used to route traffic, this traffic is routed from B to A.

2.4.1. Selection of AS-paths

There are several criteria an AS considers when it decides which path it prefers. These criteria only apply if the paths are *exactly* the same. Hence, if there are two AS-paths and one leads to a more specific prefix, the criteria are ignored and the path towards the more specific prefix is chosen. Additionally, ASes can also not prefer a less specific prefix over a more specific one.
The first of these criteria is 'local preference'. This criteria is an AS-specific one. This means that depending on its configuration the Border-router selects a route. Border-routers often decide which path to prefer, based on the business relationships to its neighbors. Paths learned from customers are preferred over the ones learned from peers. Paths learned from peers are preferred over the ones learned from providers. See Figure 2.5 for an example where AS A gets two routes towards the same origin, one announced by a peer and one announced by a customer.

The second criteria is the 'path-length'. If an AS can't make a decision based on local preference (e.g. because all ASes originating the conflicting paths are providers), it decides based on the AS-Path length. This means that the route with the lower number of ASes along the path to the originating IP-prefix is chosen. See Figure 2.6 for an example where AS A gets two routes towards the same origin, both announced by a peer.

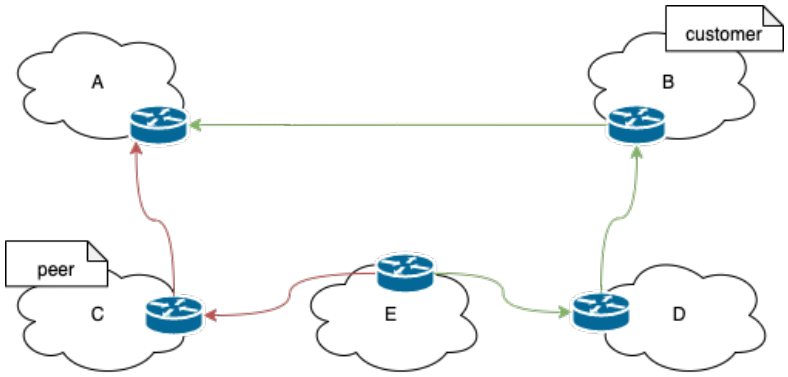

Figure 2.5.: Selection of paths that are announced by a peer and a customer:
A gets two paths towards *E*. One from its peer *C*, the second from its customer *B*. Because of the criteria 'local preference', *A* selects the path via *B*.

If the second criteria still does not break the tie, the AS chooses a path based on AS-internal routing policies. Every Border-router maintains a list of criteria which are used to make a decision. The further down this list goes, the more 'random' the compared values get, just to break the tie. Therefore there will always be a criteria which finally breaks the tie. We will not discuss these AS-internal routing policies further in this work as we are focusing on exterior routing between ASes and not on AS internal routing policies.

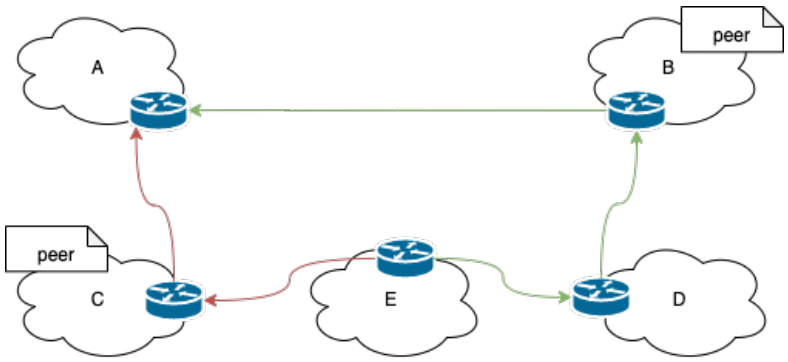

Figure 2.6.: Selection of paths that are announced by two peers:
A gets two paths towards *E*. One from its peer *C*, the second from its peer *B*. Because of the criteria 'path length', *A* selects the path via *C*.

2.4.2. BGP loop-prevention

One feature of BGP that is necessary to understand for a specific attack type later, is BGPs *loop-prevention mechanism*. As its name already tells, loop-prevention prevents emergence of loops in AS paths. Consider Figure 2.7 for an example.
AS *C* announces its ASN and prefix to *D*. The path emerges and finally reaches *A*. *A* also

prepends itself to the path and sends it to its own neighbors via an UPDATE message. *C* is one of *A*s neighbors and therefore also receives the path. It (*C*) recognizes that it is already on the path and that it therefore wouldn't make sense to prepend itself to this path. Hence, the loop-prevention mechanism applies and causes *C* to drop the received AS-path.

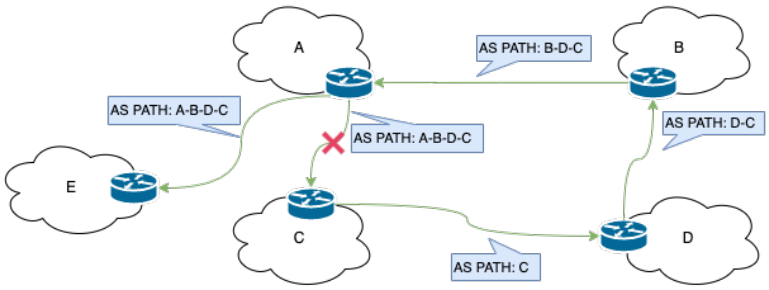

Figure 2.7.: Loop-Prevention: *C* announces a path. This path spreads via *D*, *B* and *A*. *A* passes the path on and it reaches *C*, again. *C* finds itself already on the path and drops it, to prevent emergence of a loop.

2.5. Misconfigurations of Autonomous Systems

BGP was designed in 1989. In this time there were made little attempts towards security[6]. This also implicates, that BGP-announcements are neither encrypted, nor signed. Consequently, there is no method to validate if an AS-path stems from the pretended source or not. Thus every AS in the Internet is able to announce any desired ASN and prefix. It is not validated if the announced ASN and prefix actually belong to this AS or not.
There have been several incidents in the past, where failures of ISPs or network administrators of Autonomous Systems have lead to severe impact on global routing. One of those events was the unreachability of YouTube due to a misconfiguration by Pakistan Telecom in 2008 [11]. Originally YouTube's AS was announcing its ASN and a list of prefixes. When the Pakistani government decided to block access to YouTube, things went wrong. Due to a misconfiguration, Pakistan Telecom not only started to announce one of YouTube's IP-prefixes, but also announced a more specific one than YouTube itself. Recall, that the problem with that is, that other ASes prefer routes to more specific over routes to less specific prefixes. In Figure 2.8 there is an example for this incident. AS *A* announces a list of its prefixes and its ASN. AS-paths emerge and everything works fine. However after some time passes, *D* announces one of *A*s prefixes too. As *D* announces the more specific prefix for `134.169.9.1`, other routers will adopt *D*s route and traffic will be routed towards it.
Other incidents were a traffic loss of a high percentage of worldwide traffic caused by a mistake of a turkish ISP in 2004 [12] or the unreachability of Facebook in 2021 [13]. How-

[6] In fact there exist methods to secure BGP like *BGPsec* and *RPKI*. BGPsec and RPKI both provide the option to sign AS-paths before announcing them [9, 10].

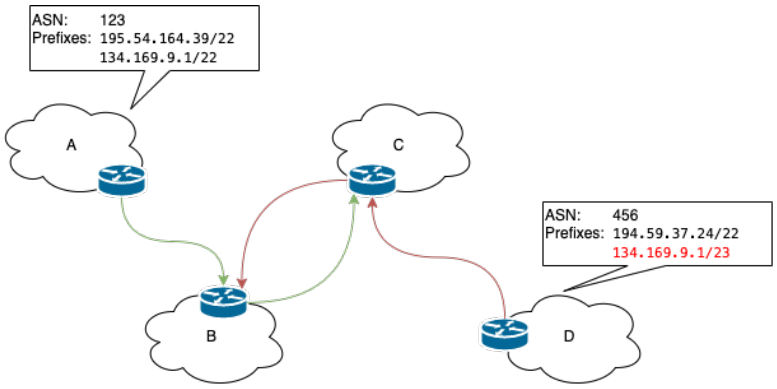

Figure 2.8.: Misconfiguration of an AS: *A* and *D* both announce the same IP-prefix. *D* announces the more specific one. Thus *B* and *C* will adopt *D*s route to `134.169.9.1`

ever Facebook's unreachability was not caused by another AS announcing the wrong prefixes, but because Facebook withdrew its routes due to internal occurrences. Several other events have also been observed. [14]

2.6. BGP-based attacks

While misconfigurations can have severe impact on routing, the absence of authentication policies can also be exploited on purpose. Attackers who are in control of an AS can cause traffic interruption for other ASes or even eavesdrop victims traffic. We will introduce the most common BGP-based attacks in the following.

2.6.1. AS-path poisoning

In AS-path poisoning an adversary poisons a victim's AS so that other ASes will stop routing traffic via the victim's AS. To do so, AS-path poisoning exploits BGPs loop-prevention mechanism. Consider an adversary, controlling AS *A* wants to poison AS *B*. To achieve this, *A* not only prepends itself with its ASN and IP-prefix to AS-paths it announces to its neighbors, but also AS *B* with its information. If this path is one that already includes *B* the path is immediately dropped by the loop prevention mechanism. If the path does not contain *B* yet, it is announced to *A*s neighbors. If it actually reaches *B* after some additional hops and *B* prepends itself to it, it will be dropped then, because of the loop-prevention mechanism. Thus, all paths that contain *B* and were also routed via *A* will be dropped and therefore *B* is not reachable via these paths anymore. The more paths containing *B* that are also routed via *A*, the less reachable *B* becomes.
It is not only possible to poison single ASes but also entire AS paths. This simply means that the adversary poisons every AS along this path.

2.6.2. BGP hijacking attacks

BGP hijacking attacks work similar to the previously introduced misconfiguration. (See Section 2.5.) If an attacker wants to eavesdrop traffic, destined to a certain prefix, it can pretend to be the legitimate owner of this prefix by announcing it via its own AS. It can even announce a more specific IP-prefix than the legitimate owner. Therefore traffic destined to the victims AS will now potentially reach the adversary instead of the victim. The attacker now is in control of the hijacked data and can use it for its own purposes. After that, the traffic is dropped. Consider Figure 2.9 for an example: AS A, the legitimate owner of prefix `134.169.9.1/22` announces this prefix and its ASN. This valid path (represented by the green arrows) spreads via B and reaches C. D, the adversary also announces As subnetwork, but with a more specific prefix (represented by the red path). It is very likely to happen, that C will drop the valid path to A it got from B and rather adopts the bogus path towards `134.169.9.1/23` it gets from D.

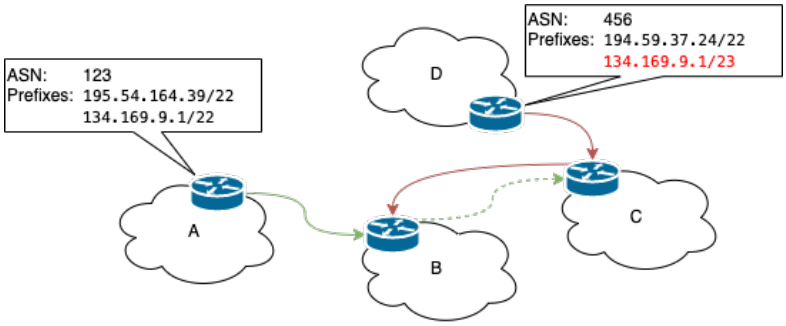

Figure 2.9.: D hijacks As prefix `134.169.9.1/22`, by announcing a more specific one, `134.169.9.1/23`. Traffic that was originally destined to A will now reach D and is therefore *hijacked*.

2.6.3. BGP interception attacks

The problem that comes with BGP hijacking attacks is, that the victim is very likely to notice that its traffic has been hijacked. This is, because the adversary drops the hijacked traffic. The problem of dropped data is addressed by *BGP interception attacks*, a more specific kind of BGP hijacking. An adversary precedes as for a BGP hijack. But instead of dropping the hijacked traffic it tries to pass it on to the AS it was originally sent to by the victim. If the adversary succeeds, the victim will possibly not notice that its traffic has been intercepted because it actually reaches its intended destination.

To be able to pass on the hijacked traffic to its original destination A, the adversary D must be in control of a multi homed AS with at least two Border-routers $R1$ and $R2$. Via $R1$ the hijacked prefix is announced. $R2$ is used to route the hijacked traffic towards the original destination, the victims AS A. If the adversary only had one router $R1'$ to announce the hijacked prefix and to send out the hijacked traffic towards the victim, the hijacked traffic would always be routed back to the adversary. This is, because neighbored routers of

$R1'$ assume that $R1'$ is the valid destination of the victims traffic. Therefore two Border-routers are needed to run an interception attack. Consider Figure 2.10 for an example. The red AS-paths represent the bogus AS-path announced by the adversary D. The green AS-paths represent the valid AS-paths announced by A. The green dotted AS-path from B to C represents that B is likely to prefer the bogus path over the valid one, as it contains the more specific prefix.

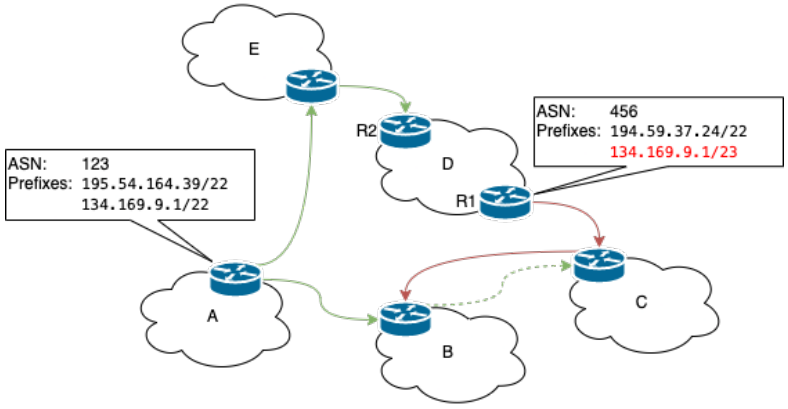

Figure 2.10.: D hijacks As prefix `134.169.9.1/22`, by announcing a more specific one, `134.169.9.1/23` via R1. Traffic that was originally destined to A will now reach D and is therefore *hijacked*. On top of that D passes on the hijacked traffic via R2 which did not learn the bogus AS-path announced via R1. Therefore E is used to pass on the traffic towards the victim, A.

The greatest challenge for the adversary that comes with interception attacks is, to always preserve a valid AS-path to the victim. If it fails to do so it is unable to route the hijacked traffic on to the original destination, the attack fails and the victim is likely to take notice of the attack.

3. Related Work

There already exists a broad body of research regarding AS-level interconnection, routing topology and its visualization on AS-level. Donnet et al. give an overview on early research towards mapping Internet topology [15]. Before this is possible, ASes have to be matched to their geo-location. Winter et al. outline why this problem is a non-trivial task and how it can be tackled [16]. Our work is strongly related to Acharya et al's study [4] which maps Internet topology and analyzes which ASes are important hubs in the Internet Backbone. Beside them there exist several approaches which all model Internet topology. As our work is motivated by security- and privacy-related questions concerning BGP we will also cover BGP-related attacks in this section. Aforementioned topics split the related work section into four groups. The first group of work takes a general look on how traffic is routed on AS-level and how routes are chosen by ASes. The second part solely tackles the task of modeling Internet topology and the third group of work builds on those (either self created or already existing) models and investigates into further questions. Finally we take a look at BGP-based attacks. This thesis resides between those groups as we are motivated by security- and privacy-related questions. Furthermore one of our main contributions is a methodology which analyses topology and further steps perform evaluation based on this methodology.

3.1. AS-level routing

Facing the task of analyzing Internet topology it stands to question, why routers choose exactly those paths they do. Early research of Gao et al. identifies the hierarchical structure of AS interconnection [17]. They generally split relationships of two ASes which are connected with each other into three groups. An AS can either be a peer, a customer or a provider for another AS. This means that two ASes either have a relationship on equal terms (which mainly concerns two ASes of the same size) or charge each other for forwarding the other ASes traffic. This model of how ASes choose routes follows the *valley-free-property*. This means that an AS-path like *A*-*B*-*C* with AS A and C as providers to B is unlikely to exist. Our work builds on their theory, as we have to understand why ASes choose exactly the paths they do, for our analysis. Following Gao et al.'s theory, it is also possible to infer relationships of two connected ASes just from analyzing their routing tables [18]. This method of inferring relationships can be used to hypothesize on AS-paths between two given sets of origin and destination ASes.

3.2. Mapping Projects

There are generally two methods when analyzing which route traffic takes. The first method is based on tools like *traceroute* or *tracepath*. These tools provide an IP-level view on how data is routed on the way to its destination. When using traceroute, IP-addresses have to be resolved to their corresponding ASN. This is less accurate than directly using BGP-data for analysis (e.g. projects like RouteViews [19]). Another flaw of using active measurement methods like the traceroute/tracepath approach is that those measurements are also dependant on internal routing protocols (e.g. RIP) and their policies. Dependant on the point of measurement in a network, routing might be different when running a traceroute at another point in the network. In contrast to that using BGP data does not have this flaw because BGPs routing does not depend on AS-internal routing policies. From these two methods we use the method of passive measurement in our work. This brings the benefit, that we do not have to actively measure AS-paths, but are able to draw on existing data. Claffy et al. introduce CAIDAs[1] tool Archipelago (Ark) [20] which is an Internet measurement platform and provides AS-level topology data [21]. Ark uses multiple data sources and processes collected data. One of their main sources is the University of Oregon's *Route Views Project* [19]. RouteViews is connected with over 600 peers all over the world who collect BGP data. Ark also provides data collected by *RIPE*[2] Their third data source are their own monitors which are similar to RIPe's collectors and also collect BGP data. Ark uses both methods of measurement, IP- and BGP-level data. Equally to Ark, we also use RouteViews as our data source for AS-paths. Therefore this leaves us with the task of of filtering ASes and the step of analyzing data. We originally planned on using Ark's data for our work, but it turned out that the project drew to an end before we were able to request their data. Another approach is given by Winter who also builds a representation of Internet topology [22], based on data by UCLA's Internet Research lab [23] and RouteViews [19]. Their data is only BGP-based. Durairajan et al. present their study *Internet Atlas* which differs from aforementioned work as it aims to model a combination of the *Physical Internet* (e.g. data-centers, ISPs facilities, IXPs etc.) and dynamic data such as AS-paths [24]. In contrast to their study we concentrate on mapping only AS-paths and the physical location of ASes. We do not take into account other infrastructure like Durairajan et al. do.

3.3. Topology Based Studies

Acharya et al. investigate which ASes around the Internet are important for global routing and how many of them are based in censorious countries. They find out that a third of those ASes which they identify as *Key Players* lies in censorious countries. Furthermore

[1] CAIDA stands for 'Center for Applied Internet Data Analysis' and is located at the University of California's San Diego Supercomputer Center

[2] RIPE is a Regional Internet Registry (RIR) providing global Internet resources.

those Key Players are contained in 20% of the paths they consider in their tests. This leaves the conclusion that censorious countries filter a substantial portion of the worlds Internet traffic. Our work is strongly related to their study. Despite the fact that two of our contributions replicate parts of their work, we are curious to ask if their findings can be reproduced four years later. Over this time the number of ASes in the Internet has almost doubled [4, 5] and this growth might have also brought changes in topology. Karlin et al. study the impact different countries have on worldwide routing and identify which countries are *Big Players* in the global Internet [25]. They find out that the United States are the most important country in global routing, followed by Great Britain and Germany whereas censoric countries are less important hubs for routing. In contrast to our work, Karlin et al. focus on analysis of countries as hubs. Our work takes another perspective as we consider ASes as points of control and not only entire countries. Edmundson et al. do similar work but use measurements on IP-level from RIPE Atlas *probes*[3] to study hegemony of certain states in global routing [26]. One of their key contributions is to identify *transnational routing detours*. An AS-path contains a routing detour if its origin and its destination are in the same country but among the path are one or more ASes which are not located in this country. They find out that routing detours exist in 85% of their studied paths. Strategically choosing peering nodes for routing traffic can reduce routing detours to 38% of studied paths. Based on their findings they also create *RAN*, a tool which enables users to circumvent certain countries when surfing the Internet [27]. Another approach that goes in the same direction is given by Shah et al [2]. What differs their work from Edmundson et al.'s is that they use BGP routing tables and not actively measured data like traceroute/tracepath. They find out that out of 7 billion paths they take into account 2 million contain routing detours. Similar to Edmundson et al.'s tool RAN, we analyze if there are ASes which cannot be circumvented when routing traffic. However their work differs in two points as they investigate on IP-level and focus on countries and not on single ASes, like we do. Next to analysis of global Internet topology there also exists work which observes topology of specific countries. Roberts et al. map connectivity of several nations as India, Russia, Sweden and more and compare their topologies [28]. As to expect, they find out that censorious countries have much fewer points of control than others. This makes it much easier for governments to control those countries traffic or even shut it down entirely. For example Iran and Libya have only a single AS which connects them to the outside Internet. In contrast to that, European countries are much more interconnected and make it much harder for governments or single organizations to control a countries Internet traffic. In contrast to their work, we not only consider ASes inside a specific country as potential points of control for this country. In fact we also consider ASes outside a certain country as points of control for this country. Next to analysis of worldwide routing there also exists work concerning single countries. Wählisch et al. study topology of the German Internet [29]. They analyze business sectors (e.g.

[3] RIPE Atlas is a Internet measurement platform and runs a network of devices called *probes* which measure Internet connectivity.

government sector, medical sector, etc.) and how they peer inside of Germany. One interesting point they line out is, that services that can be associated with the governmental sector mainly use two ASes to peer with. The owner of those ASes are *Deutsche Telekom* and *Versatel*. Similar to their work, Oh et al. take a closer look on Korea's Internet [30]. After inspecting a topology graph of the Korean Internet they conclude that ASes among Korea are stronger connected than the worldwide Internet graph and are therefore more efficient in routing than the Global Internet. Zhou et al. do analysis on how the Chinese Web is build [31]. Even though development of the Chinese Internet is more influenced by central planning than by market-based competition its internal topology is similar to global topology. This is a finding we can also confirm, as we consider China as one country of interest in our analysis. Another study that goes into a completely different area of the world is given by Fanou et al [32]. They investigate into the interconnection of ASes based in Africa. One remarkable observation they make is, that Africa has a lack of interconnection between their ISPs. This leads to the usage of ISPs outside of Africa to interconnect the continent. This fact also leads to high costs for routing traffic. Their findings match with our insights as we note a general underrepresentation of ASes among African countries.
Aforementioned work focuses either on a global (international) or a local (national) view of the Internet. In contrast to that we take both views into account and compare if local hubs are equal to global hubs. Furthermore we analyze if there are certain ASes which cannot be circumvented when routing traffic.

3.4. BGP Security

Early steps of research take a general look on which properties of BGP can be exploited in which way. One of these studies is by Nordström et al. They investigate into several attacks which exploit flaws in BGPs design [33]. We outline some of them in Section 2.6. They also discuss existing countermeasures and show, why those countermeasures are partly ineffective and partly too sophisticated in deployment. Our work confirms their conclusion, that there is still a lot of work to be done until misconfigurations or attacks in routing can be effectively prohibited. Further work investigates into the specific attack types. Hijacking- and interception-attacks are two types drawing a lot of attention. We introduce these two attack types in our study and also analyze a setting in which a hijacking attack is performed by Pakistan Telekom (intentionally or not). Birge-Lee et al. investigate into this topic and come up with a new approach to make these attacks even more feasible [34]. Sun et al. show which impacts attacks on BGP can have. They propose a method to uncover Tor-Users which is enabled by BGP-interception attacks [3]. Their finding motivates for our work, as we find out that there are few ASes which cover a high amount of AS-paths in the Internet. Therefore interception attacks become less important for attacks like unraveling Tor-Users if data gotten from the few hubs is already sufficient for that.

4. Methodology

This section describes our method of analyzing AS-paths. Our method calculates several metadata parameters, inferred from collected AS-paths. We then describe how we take a local and a global perspective on the calculated metadata. Additionally we explain how maps that correspond to this metadata are created. Except for one parameter, the following steps are the same for both, the local and global perspective of evaluation.

4.1. Analysis of AS-paths

Our methodology of analyzing AS-paths consists of several steps. These steps are as follows:

1. Gathering AS-paths
2. Filtering AS-paths
3. Localization of ASes:
   - Matching IP-prefixes to ASes among paths
   - Localizing IP-prefixes extracted in previous step
4. Calculating visualization-metadata
   - Obtaining the path-ratio
   - Obtaining the hub-score
   - Obtaining the essentiality-score
5. Mapping ASes by the geolocation of their IP-prefixes and visualizing AS-paths

4.1.1. Gathering AS-paths

The first step towards analysis of AS-paths is collecting them.
Our data is extracted from publicly available RIBs of several Border-routers. Obviously this method can only bring a reliable set of data if there is a broad spectrum of measurement-locations. Measurements from single locations can only show topology of their closest neighbors. In the following we use the term *dataset* for the collected AS-paths and all ASes along the paths.

4.1.2. Filtering AS-paths

After collection of AS-paths and ASes along these paths, data can be filtered for specific path-origins, path-destinations and time-ranges. This step is needed for taking only paths

between specific ASes and in specific time-ranges into account. Depending on which perspective of analysis is taken, different filters have to be applied in this step. This is the only step where analysis differs for the global and the local perspective. If we use the term 'filter/filtering' in later steps we refer to this step of filtering AS-paths.

4.1.3. Localization of ASes

We proceed with matching ASes along the paths to their geographical locations. Despite the fact that *Regional Internet Registries (RIR)*[1] like RIPE, APNIC, etc. provide information about registration countries for every AS, this data is only legally relevant and does not represent the actual router locations of ASes. Hence, ASes have to be located by another method. This thesis uses the method of first analyzing, which IP-subnetworks are maintained by every AS and then using a *Geo-IP service*[2] to locate the investigated subnetworks. One challenge that comes with locating IP-subnetworks is, that the location of ASes can be ambiguous. Consider an AS X which announces its network Y with an IP-prefix /24. Looking up this prefix in a Geo-IP database returns a more specific prefix /25 of Y. This prefix locates to a location in Germany. The remainder of Y/24 locates to France. It is unclear whether this AS should be located to Germany or France. One way to decide, is to analyze all possible locations for an AS and then taking the country which is the most referred one among all prefixes. This approach bears the obvious flaw of discarding information. As this thesis aims to provide security and privacy-related information, another method is chosen here. Instead of only localizing an AS to one location we take all locations into account. How we preserve a clear look in our visualizations later on, will be discussed in Chapter 5.

4.1.4. Calculating visualization-metadata

To measure importance and essentiality of ASes, these parameters are quantified. In the following we describe some metadata parameters that are measured for each AS, in order to do so.

Obtaining the path-ratio

The first of these parameters is the *path-ratio*. It describes the percentage of paths, each AS appears on where it is *not* the first or the last element (and therefore the destination or origin of the AS-path). We choose to not take origin and destination hop into account, because doing so would distort analysis. Consider an AS that is highly present as origin or destination but does not carry any traffic for other ASes. This AS would have a high path-ratio even though it does not play a role for routing of other ASes traffic. See Figure 4.1 for

[1] Regional Internet Registries are responsible for assignment of IP-address space, ASNs, etc. in their respective regions.

[2] Geo-IP services provide data about Autonomous Systems, their related IP-subnetworks and geographical locations of IP-subnetworks.

an example of evaluating path-ratio. Pseudocode for evaluation of path-ratio is provided in Algorithm 1.

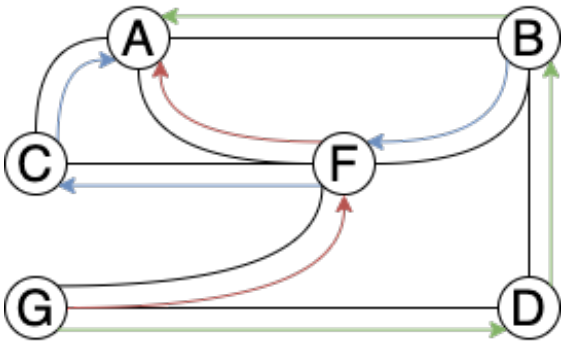

Figure 4.1.: In this example F has the highest path-ratio. It is along the red and the blue path. Therefore its path-ratio is 2. G has an path-ratio of 0 as the red and the green path start in it but it is not traversed by any path. For the same reason, A's path-ratio is also 0. D's, B's and C's path-ratio is 1 as they are each traversed by one path.

Algorithm 1 Evaluation of path-ratio

```
1: procedure GLOBAL-PATH-RATIO(asPaths, autonomousSystems)
2:     for as in autonomousSystems do
3:         pathRatio[as] ← 0
4:         for path in asPaths do
5:             if as in path and as ≠ origin(path) and as ≠ destination(path) then
6:                 pathRatio[as] ← pathRatio[as] + 1
7:     return
```

Obtaining the hub-score

For a second parameter, we count, how many ASes are needed to intercept 90% of paths among the dataset. This follows the approach of Acharya et al. [4]. We proceed as follows:

1. Rank ASes among the dataset by their path-ratio
2. Choose the highest ranked AS, remove it from the rank and add it to a list of hubs
3. Evaluate on how many paths among the dataset this AS is present (As for calculation of path-ratio, start/origin and end/destination of path do **not** count)

We do this process repeatedly, until the chosen hubs cover at least 90% of paths among the dataset. We call the number of hubs *hub-score*.

Obtaining the essentiality-score

A last parameter that is evaluated for each AS is the *essentiality-score*. This parameter is crucial for evaluation of essential ASes.
Identification of essential ASes first brings up the task of defining the term 'essential'. One option is that an essential AS is such one that is relevant for every single AS-path in our dataset. Removing it would result in disruption of two ASes that were formerly connected by an AS-path. Another option is, that an essential AS is such one that must be contained in every single AS-path in our dataset. Removing it would cause a disruption of every pair of ASes that were formerly connected.
The first option would result in a huge set of essential ASes, because disruption of a single connection between two ASes would make an AS essential. Latter of the methods would result in an extremely small, possibly empty set of essential ASes, as every single path had to be dependant of the removed AS. Therefore we choose to quantify essentiality of ASes and always relate it to a specific AS. This means that AS X's essentiality for AS Y is increased for every AS that gets disconnected from Y when removing X (and therefore every path that contains it) from our dataset. Note that an AS's rank is not increased for the trivial case of removing itself and disrupting its own connection to another AS. The higher a rank of an AS for another AS gets the more essential it is for it. Consider Figure 4.2 for an example of calculating the essentiality-score.

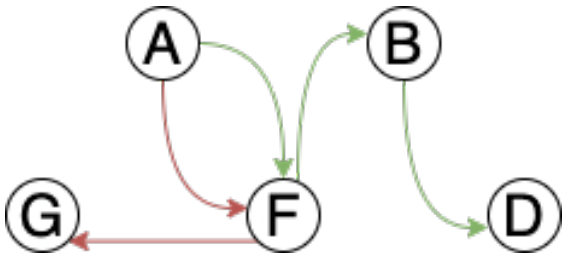

Figure 4.2.: In this example F has an essentiality-rank of 3 for A as G, B and D get disconnected from A when F is removed. B only has an essentiality-rank of 1 for A as only D gets disconnected from A when removing it. AS G has an essentiality-rank of 0 for A because removing it only disconnects itself from A.

4.1.5. Mapping ASes by the geolocation of their IP-prefixes and visualizing AS-paths

After data has been processed and all metadata has been calculated it remains the step of actually mapping Autonomous Systems and AS-paths.
First, all ASes contained in the dataset are mapped by the geolocation of their largest IP-prefix. ASes are represented by nodes. Node size is depending on the previously calculated path-ratio.
The second step maps all AS-paths from the given dataset. Paths are represented by di-

rected edges between Autonomous Systems.
More detailed decisions on implementation are described in Chapter 5.

4.2. Identifying Hubs of Control and essential ASes

After analysis of ASes is done it has to be evaluated which ASes act as Hubs of Control and are also essential ASes. This evaluation is done by analyzing path-ratio and essentiality-score. Approaching this task, we take two different perspectives on hubs. The first question we are curious to ask is, if there are global hubs. Our second focus are local (national) hubs for specific countries.

4.2.1. Identification of global hubs and essential ASes

In order to identify global hubs, AS-paths have to be analyzed as outlined in Section 4.1. We choose to not take all existing AS-paths into account, because of the sheer amount of data this would result in. On top of that, rankings like *Tranco*[3] show, that a high amount of traffic targets a small fraction of destinations. Based on these arguments we choose all ASes that resolve from *Tranco's Top Site Ranking* [36] for our set of origin ASes. Although we do define a set of origins we do not restrict the set of destinations, but take paths to all destination-ASes we can find. The set of origins we made out is applied as a filter in the process of analysis as outlined in Section 4.1.2. For evaluation of global hubs we take the path-ratio for each AS into account. We argue that the more occurrences an AS has on all paths, the more important it is for routing between the set of origins and destinations. The set of hubs are exactly those ASes that are under the first X ranked, concerning their path-ratio. X stands for the hub-score, evaluated in Section 4.1.4.
Next to identification of hubs, we analyze the previously calculated essentiality-score in our evaluation. We provide several graphs on how the essentiality-score for certain ASes changes over the years.

4.2.2. Identification of local hubs and essential ASes

To identify local hubs we have to choose sets of origins and destinations. As origins we take the set of ASes which resolve from the sites among the Tranco list again. For destinations we take ASes of ISPs for our country of interest. Choosing ISPs is for two reasons. The first reason is, that we want to identify Hubs of Control for one specific country. Therefore we take only ASes of this specific country into account. One could question why we do not take all paths to all ASes in this country into account, like we did for the global view, but

[3] Tranco is a traffic analysis tool, providing rankings on most popular websites. Le Pochat et al. show, that existing Popularity Rankings can be easily manipulated by adversaries [35]. As a result they provide a new kind of list that builds on multiple existing rankings and hardens them against manipulation by malicious actors.

instead only to ISPs. This is, because we want to observe topology how it is used by everyday-users. Since those users are mostly connected to the Internet via their respective ISP, we choose to only take ASes belonging to ISPs as our destinations. Again, we apply the set of path-origins and path-destinations as filters in Step 2, of analysis. We choose the same method for evaluating local hubs as presented for global hubs in Section 4.2.1.
On top of that we investigate if there are ASes from those local hubs that can not be circumvented when routing traffic. For this investigation we compare the essentiality-scores of the top ranked ASes. Beside that, we analyze if certain countries ASes have higher essentiality-scores than other countries ASes and relate our insights to the potential of Internet-censorship in those countries.

4.2.3. Analysis of events in historical BGP-data

Next to evaluation of hubs, we also want to investigate into cases where an analysis of AS-paths could tell of misconfigurations and attacks which affected worldwide routing. Some of those incidents have already been outlined in Section 2.5. In cases where rerouting of AS-paths happened, we expect to see that in visualizations.
For analysis we first find out at which times the incidents appeared and which ASes where affected. We proceed with mapping AS-paths before, during and after the incidents. The affected ASes are applied as path-origins in Step 2 of analysis. After analysis is done we compare visualizations for the time before, during and after the incident.

5. Implementation

This chapter builds on methodology of Chapter 4 and describes the practical steps of implementation that were taken to build an analysis tool for AS-paths and accomplish Contribution 2. The implemented tool is mostly written in Python, as this is the easiest way of collecting and processing BGP-data, because of an existing API for route collection.

5.1. Route Collection and Extraction of ASes

The main source for BGP-data that is used in this thesis is the *Routeviews Project* which is maintained by the University of Oregon. It "was originally conceived as a tool for Internet operators to obtain real-time BGP information about the global routing system from the perspectives of several different backbones and locations around the Internet" [19]. To collect data, they and their partners deployed several *Collectors* all around the globe and capture BGP-announcements they receive from their neighbored ASes. This means that Collectors act like BGP-routers and receive BGP-announcements from their neighbors but do not pass them on to their own neighbors. The collected data is provided via an archive that lists the different Collectors. To access their archive, CAIDA provides a C library called *libBGPStream*. This library provides easy access to RouteViews data and gives several options to filter it. They also provide a Python package, *PyBGPStream* that builds on libBGPStream and allows direct data access from Python scripts. Our first step is, to use PyBGPStream to access RouteViews. We make use of PyBGPStreams filter functions to filter routes for destination and origin that have been given as input parameters by the user. Doing that we make sure that we only take relevant routes into account and also do not have to do this step later on. We also filter by start and end time. This filter is mainly needed because not applying it would result in a data-set of enormous size. It simply does not make sense to not filter data by a time-range. This also brings the benefit that during evaluation, routes from different time ranges can be compared and it can be analyzed how BGP-data changes over time. After route-collection we remove duplicates and make sure we only have unique routes. One could question why there are duplicates around the collected data. This is because RouteViews relies on a set of Collectors and multiple Collectors may get equal routes from their neighbored ASes. The next step takes the collected routes and extracts all ASes from them. Autonomous Systems and AS-paths are now ready for further processing which is described in the following.

5.2. Matching of Autonomous System Numbers to IP-prefixes and Localization of IP-prefixes

After collection, Autonomous Systems are resolved to their geographical location. We use *MaxMind*s [37] free Geo-location databases to tackle this task. First ASNs have to be attributed with their IP-subnetworks. In a second step we match IP-subnetworks with their geographical location. This is, because MaxMind does not directly assign locations to a given ASN but provides several sets of information. The most important ones for us are one set which assigns IP-ranges to ASN and another one which assigns geographical locations to IP-ranges. By joining these two sets on the shared attribute of IP-subnetworks we can indirectly locate a given ASN. Since one AS can be assigned to multiple IP-subnetworks, it might also have multiple geographical locations. To not lose information, we store a list of locations for every Autonomous system. Locating ASes to the locations of all their IP-subnetworks would result in a enormously messy map. Thus, we decided to only locate ASes to the location of their largest IP-subnetwork. In the interactive visualization one can see all other locations of an AS when hovering on it.

5.3. Visualization of Routes

The last step is to visualize the processed ASes and paths. We use the *plotly* [38] Python library to visualize a map with nodes and edges. Nodes represent ASes and paths are represented by one or more edges. Before we provide node and edge data to plotly, we evaluate the size of ASes dependent on their path-ratio. By doing this we highlight highly used ASes by increasing the size of their corresponding node. After this last step of process, we pass nodes and edges to plotly and finish with the step of actual mapping.

5.4. Adjustments for analysis of local hubs

One parameter that has to be changed when analyzing local hubs, is the geolocation of ASes which are located in multiple countries. Recall that we located ASes to the location of their biggest IP-subnetwork in previous steps. As we are interested in those portions of an AS which are located in our country of interest, we locate the ISPs Autonomous Systems to this country, even if a bigger IP-subnetwork is located in another country.

5.5. Adjustments for analysis of historical BGP-data

As outlined in Section 2.5, misconfigurations are often caused by the wrong announcement of specific IP-prefixes. Thus, not the location of an entire AS, but the originating AS (and consequently also the location) of a single IP-prefix changes.

Locating AS-paths by their ASN would still result in the same visualization as before an incident. Hence, for this step we need to located affected IP-prefixes to their originating ASes.
Conveniently, data collected by RouteViews not only provides AS-paths but also the IP-prefix of the origin AS. Therefore we can filter paths by IP-prefix as path-origin. Consequently, if the related AS of an IP-prefix changes, this is observable in data.

6. Challenges and Limitations

There are several challenges when analyzing Internet topology. Some of them stem from the fact, that there are multiple different data sources needed. Other challenges are caused by an extremely high amount of data. This section describes and discusses major challenges we ran into. We will also present, how we tackled these challenges and/or which restrictions we had to make.

6.1. Geographical data

To be able to locate Autonomous Systems, geographical data is needed. As mentioned before, we use MaxMinds free geodata for locating ASes. They provide their data in multiple csv-files with up to 3.5 million lines for locations of IP-networks. As the data is not provided in a single file and there is no method to directly look up the geographical location of an Autonomous System, the different files have to be connected. In a first step we decided to put the data into a database to make handling easier. The resulting tables are equal to the files that came from MaxMind. There remain two options to connect a given ASN to the location of a network. The first option is to join three tables (Autonomous Systems, Network Locations and Cities. This is, because network locations do not have coordinates as an attribute but have to be joined with cities, which are provided with coordinates.) This would result in a temporary table containing roughly 0.2 quintillion rows, which is not practically at this step. Another method is to look up data from the database step by step. This means that we first look up all subnetworks of a given AS from one table. In a next step we get the locations of those subnetworks. Finally we get the coordinates of the cities that are referenced in the location of each subnetwork from a third table. This method bears the flaw, that locating one AS takes a relatively long time. Already the origin ASes in our evaluations have around 120 subnetworks per AS and we measured a time of 52 seconds for an AS with 120 subnetworks. This does not seem to be a very long time, but given the fact that an average dataset in our analysis contains around 700 ASes it becomes clear that this problem shouldn't be underestimated. We chose to go with the latter method. To minimize runtime for analysis we cache ASes we located ones, so the runtime of our tool drastically reduces for later analyses where the same ASes are geolocated again.

6.2. Historical website rankings

As already described, we use Tranco's website ranking for the AS-path origins in our evaluation. Tranco is a research project presented in 2019. Their data goes back to December 2018. This means, that we can not rely on Tranco lists for analysis of historical data, simply because lists do not exist for this time. We decided to switch to Alexa [39] website rankings

for analysis of historical data. Alexa rankings are also part of the Tranco lists, but differ in some details. As we only had Alexa rankings going back to 2009 at hand, we had to rely on the first available list from 2009, for analysis of data, earlier than 2009. Therefore, analysis of historical data might be inaccurate in some cases.

6.3. Research of representative ASes

Originally, we planned on choosing a set of representative ASes as AS-path destinations for analysis. We planned on choosing this way to get a more overseeable set of AS-paths and keep the resulting map clean. In the process of making out representative countries for each continent and representative ASes for each country, we stumbled upon this problem. Finding representative ASes, e.g. ISPs (because ISPs carry the highest fraction of traffic for 'everyday Internet users') is a non-trivial task. Firstly, it is hard to find reliable data for ISPs in countries like Ethiopia, Egypt and countries of South-America and secondly, some countries like the United States or Australia have such a wide range of ISPs that making out the most important ones is challenging. Another reason for giving up this method is, that there is nearly no available data for ranking ISPs in their importance. On these grounds we decided to not follow this method for analysis with a global focus. Although, we were able to find out ISPs for a hand full of countries, which are analyzed in our local analysis. Investigating into ISPs for these few ASes already took a high effort. Consequently, this confirmed us in our decision to not collect ISPs for a wider range of countries.

6.4. Multi located ASes

As outlined earlier in this work, another challenge is, that a high amount of ASes cannot be located to a single spot but resides in multiple locations. This is, because multiple IP-subnetworks can be assigned to a single ASN. We came up with three ideas to tackle this challenge. The first option is to simply locate Autonomous Systems to the location of this network which has the highest amount of IP-addresses. Following this way would result in a massive loss of data. Another approach is to locate an AS to every of the locations of its IP-subnetworks. This way would result in a very confusing visualization, since some ASes have more than 5000 IP-subnetworks and are scattered all over the world. We decided to locate an AS to every of the locations of its IP-subnetworks and visualize it at the location of the biggest subnetwork (the one that contains the most IP-prefixes). The remaining locations are attached to the ASes marker on the map and can be shown as a list when hovering on it, in the interactive version of our mapping tool.

6.5. Availability of AS-path data

Attempting to analyze a representative fraction of the Internet bears the obvious challenge of gathering a representative fraction of AS-paths.
The greatest challenge for finding an appropriate dataset is, that the majority of ASes among the Internet does not route traffic for other ASes. Hence, they do not prepend themselves to AS-paths they get from their neighbors and pass them on. Therefore paths that are captured by RouteViews collectors will not receive any paths from those ASes. Consequently, we do not know how these ASes route traffic towards the origin ASes from our dataset.
Another problem is, that data-sources, especially for historical data are rare. Some organizations 'owning' ASes provide views on AS-paths contained in their RIBs. In nearly all cases one is only able to search for specific path-origins. Another challenge is, that those paths are only available via a web interface, so that there is no practical way of retrieving them automatically. On top of that, most web interfaces that provide AS-path data have a maximum of possible requests. Consequently we decided to limit our data-source to RouteViews data. There are approaches which infer unknown paths from business-relationships of ASes [4], but it is unclear whether inferred data is as reliable as already existing.
All in all these restrictions leave us with datasets which contain 0.8% to 1% of the worlds ASes. We will further consider this limitation in our later discussion.

7. Results

This section presents results we get for different perspectives of analysis. Results are based on our method, presented in Chapter 4. Findings for presented results will be discussed in Chapter 8, later. We split this section into the following subsections:

1. Analyzing AS-paths from a global perspective
2. Analyzing AS-paths from a local perspective of chosen countries
3. Analyzing AS-paths for misconfigurations of ASes

We analyze AS-paths from multiple datasets and compare results. Datasets that are used for analysis and their specific filter-parameters, applied in Step 2 of analysis are presented in a table at the beginning of every subsection. As proposed in Chapter 4 we use top pages of a Tranco ranking as the origins of AS-paths since they are the most frequented traffic destinations. To validate that our results are not dependent on the specific origin-pages contained in a Tranco list, we always perform two analyses, based on datasets with Tranco top-10 and Tranco top-20 lists as path origins.
In Chapter 4 we proposed some parameters that are calculated in the process of analysis. If we use the term "rank/ranking" over the following section, we always speak of ranking ASes by their path-ratio (descending).

7.1. Analyzing AS-paths from a global perspective

Our first steps of analysis cover latest AS-paths from October 2021. Next to observations we make for those datasets, we will also evaluate which parameters change over the past 20 years. Therefore we analyze AS-paths for the years from 2000 until today and compare the results we get.

7.1.1. Analysis of latest AS-paths

Datasets used for analysis of global AS-paths are presented in Table 7.1.
For dataset 2021_oct_top20 we ranked ASes by their path-ratio calculated during analysis. We observe, that together the top-10 ranked ASes are present on 72.6% of paths among the dataset. For the top-20 ranked ASes already 81.9% of ASes among the dataset are covered. The calculated hub-score for this dataset is 65. Hence, 65 ASes are needed to cover 90% of paths among the dataset. This is a very interesting result, as 65 ASes are only a fraction of 9% of ASes among the dataset. Another observation is, that there is a small amount of 11 ASes with a high path-ratio (higher than 5%) and the rest of ASes is only present on very few AS-paths.

dataset	path origins	path destinations	collection date
2021_oct_top10	Tranco top-10 pages	⋆	2021-10-04
2021_oct_top20	Tranco top-20 pages	⋆	2021-10-04

Table 7.1.: Datasets of latest AS-paths used in analysis with a global focus. (The ⋆ means that data was not filtered for the argument of the corresponding field. E.g. ⋆ in 'path destinations' means, that data was not filtered for certain path destinations, but paths to all existing destinations are taken into account.)

The highest ranked AS from this dataset is AS4134 which is located in China. It has a path-ratio of 25.0% and is therefore found on 25.0% of paths. The second ranked Autonomous System is AS1299, which is located in Sweden and has a path-ratio of 17.2%. Another observing is, that there are 10 ASes which can be located to the United States under the top-20 ASes. One would expect that all major ASes from our ranking are also tier-1 ASes. It stands out that especially the highest ranked AS4134 is a tier-2 AS. In total 50% of the top-20 ranked ASes are not tier-1 ASes. A detailed ranking of top 20 ASes from 2021_oct_top20 and their corresponding countries and tiers is shown in Table 7.2.

Rank	AS	Country	Intercepted paths	tier
1	4134	China	25.0%	tier-2
2	1299	Sweden	17.22%	tier-1
3	174	United States	12.68%	tier-1
4	3356	United States	13.26%	tier-1
5	6453	United States	8.63%	tier-1
6	2914	United States	8.57%	tier-1
7	4811	China	6.65%	tier-2
8	3257	United States	5.92%	tier-1
9	4837	China	5.66%	tier-2
10	6762	Argentina	5.34%	tier-1
11	8075	United States	4.89%	tier-2
12	6939	United States	4.86%	tier-2
13	12956	United States	4.36%	tier-1
14	3491	United States	4.29%	tier-1
15	58466	China	3.78%	tier-2
16	6461	United States	2.74%	tier-1
17	23724	China	2.59%	tier-2
18	4637	Japan	2.58%	tier-2
19	38283	China	2.42%	tier-2
20	4755	India	2.34%	tier-2

Table 7.2.: Top 20 ASes and their corresponding countries in October 2021

Another result are the countries, hubs are located in. A pie-chart, showing which countries host most of the hubs is presented in Figure 7.1. It stands out, that together, China (23.1%) and the United States (26.2%) have nearly 50% of hubs, while the next country is Switzerland with 10.8% of hubs.

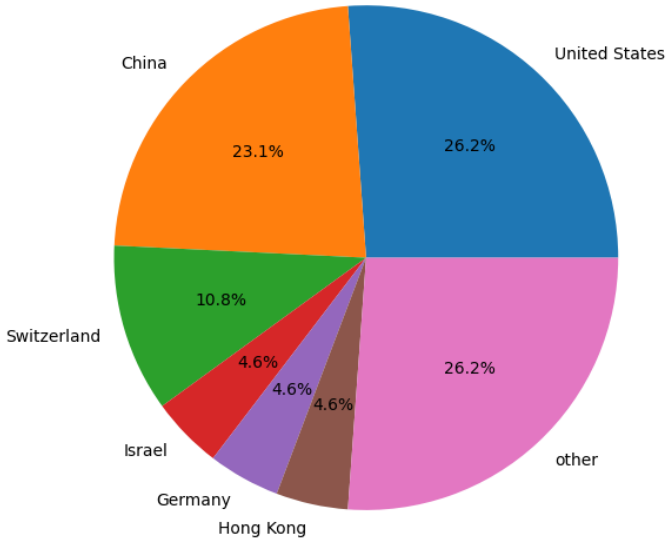

Figure 7.1.: The chart is showing countries which host most of the hubs, evaluated for global analysis.

We compare results gotten from this first dataset to results from a second one, 2021_oct_top10. Observations from this second dataset confirm our insights gotten from the first one. Again, there are around 11 ASes, being present on more than 5% of all paths. It also marks the United States as a country where many top ASes are located. Figure 7.2 compares the ranks of ASes from both datasets. One can see, that the order of ranking differs, but important ASes are the same in both datasets.

Next to data concerning the importance of each AS, we also collected data that tells of the essentiality, each Autonomous Systems has for other ASes. Calculation of the essentiality-score is outlined in Chapter 4.
For a global perspective, the essentiality-score shows, that top-ranked hubs are not necessarily also essential for routing. Recall, that every AS gets an essentiality score for every other AS that depends on it when routing traffic. We calculate the average of all essentiality-scores an AS has for other ASes. E.g. AS X has an essentiality-score of 0.3 for AS Y and an essentiality-score of 0.1 for AS Z. Hence, its average essentiality-score is (0.3 + 0.1) / 2 = 0.2. The results for all ASes that have an average essentiality-score higher than 0 are presented in Figure 7.3. What stands out is, that there are only 7 ASes (out of 700) that have an average essentiality-score higher than 0 and out of those 7, only 2 have

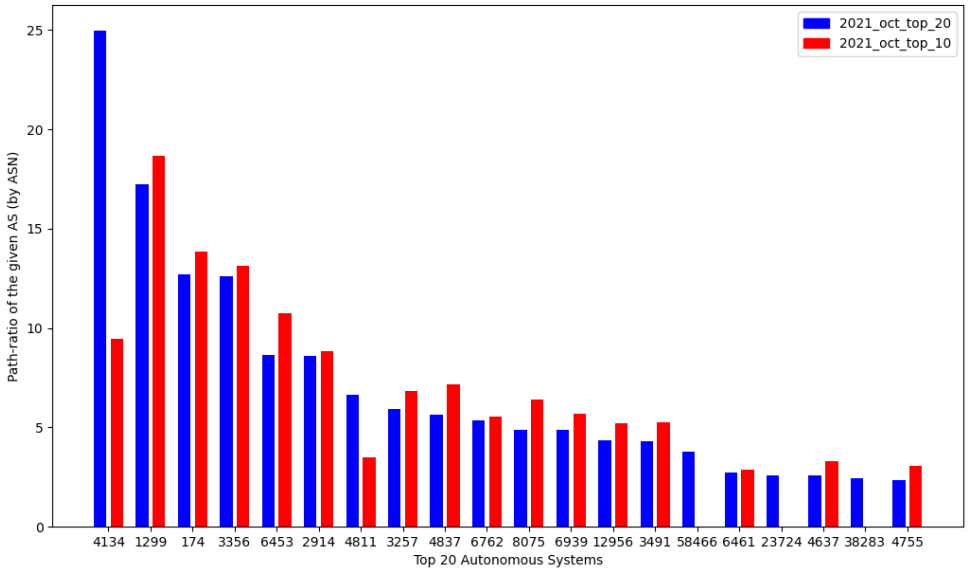

Figure 7.2.: Top 20 ASes from 2021_oct_top20 and 2021_oct_top10 compared. It shows that the ranking of ASes differs, but hubs are the same for both sets.

an average essentiality-score higher than 5%. Also these results are virtually equal for both datasets we analyzed.

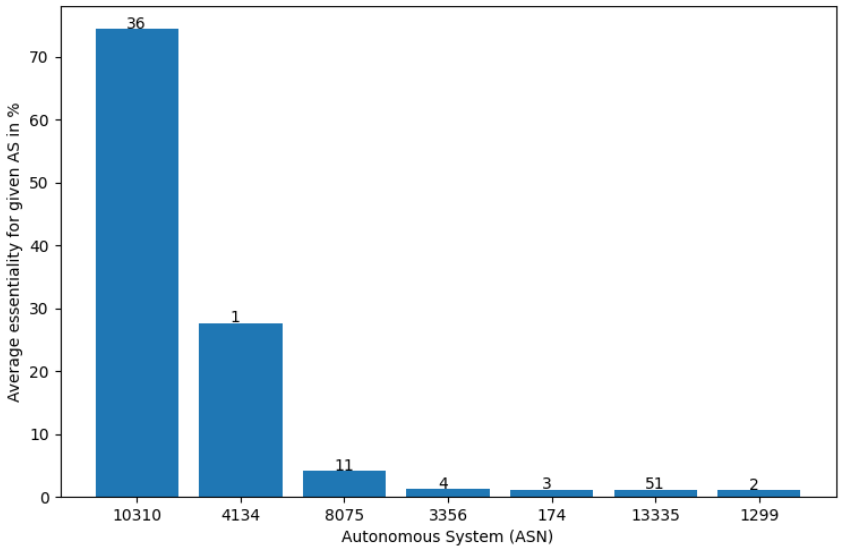

Figure 7.3.: The average essentiality-scores for ASes that have a score higher than 0 show, that essential ASes are not necessarily the ones that are also hubs. The numbers on the bars stand for the ranking the corresponding AS has in terms of path-ratio.

7.1.2. Analysis of historical AS-paths

In this step of analysis we take a look at data from 2000 until 2021 and compare how parameters changed over the past time. Datasets used for this step are presented in Table 7.3

dataset	path origins	path destinations	collection date
2021_oct_top10	Tranco top-10 pages	⋆	2020-10-04
2021_oct_top20	Tranco top-20 pages	⋆	2020-10-04
	...		
2000_oct_top10	Tranco top-10 pages	⋆	2000-10-04
2000_oct_top20	Tranco top-20 pages	⋆	2000-10-04

Table 7.3.: Datasets of historical AS-paths from 2000 to 2021 used in analysis with a global focus. (The ⋆ means that data was not filtered for the argument of the corresponding field. E.g. ⋆ in 'path destinations' means, that data was not filtered for certain path destinations.)

The first parameter we want to consider for global hubs is the hub-score of different datasets. For comparison, it makes sense to consider two variants of hub-score. While hub-score as introduced in Chapter 4 provides a good measure for total amounts of ASes, it does not state a good parameter for analysis of changing topology. This is, because it does not take the total amount of paths among a dataset into account. Consider two datasets, one of them containing 100 paths and the other 10.000 paths, both having a hub-score of 60. While for the first dataset those 60 ASes are very powerful, they are not anywhere near this power for the second one. This is the reason why we take a relative path-ratio into account for comparing datasets from different years. The relative path-ratio is simply the path-ratio divided by the total amount of ASes among its respective dataset.

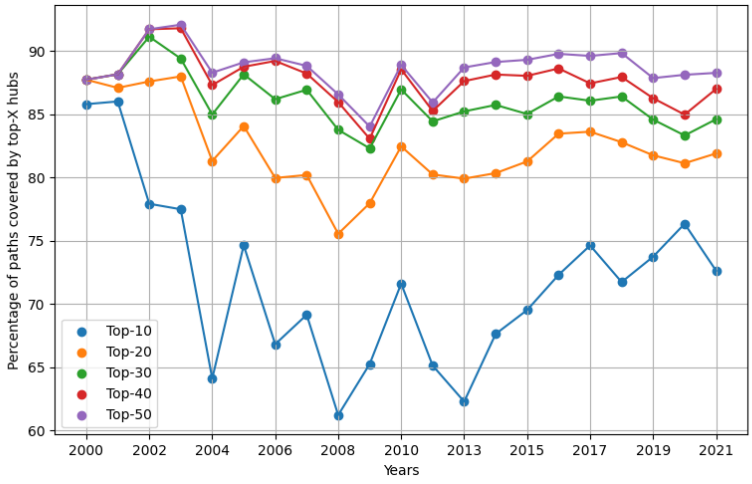

Figure 7.4.: The top-10, -20... ranked hubs (in terms of path-ratio) cover similar fractions of paths around their respective datasets over the years.

This decision is supported by an interesting observation. We plot the fraction of paths, the top-10, -20,... ranked ASes from different datasets over the years cover. See Figure 7.4. As one can see, the fractions fluctuate, but do not show a certain tendency over the years. In contrast to that, we plot the relative path-ratio for different datasets over the years. Since the number of ASes among datasets has significantly increased over the years, calculating the hub-score against the total number of ASes shows that the relative fraction of ASes in control has significantly decreased. We can clearly see this decreasing tendency in a respective plot. This result is shown in Figure 7.5.

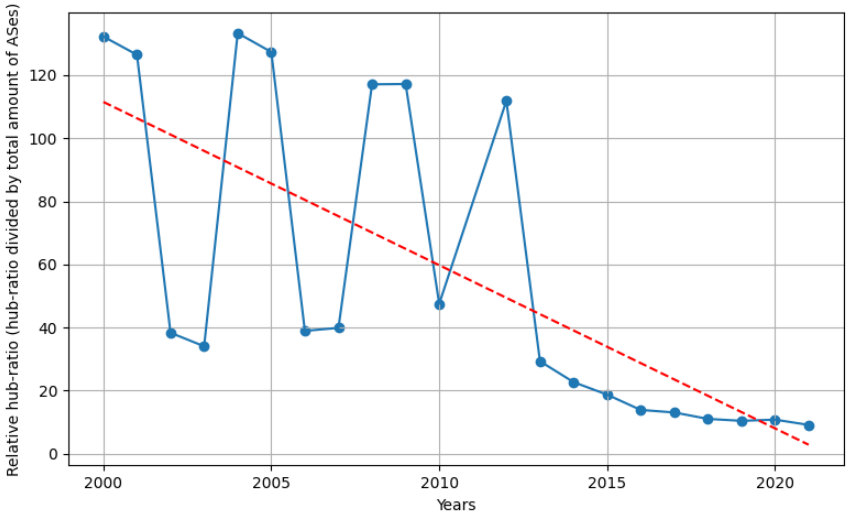

Figure 7.5.: Relative hub-score for datasets over the last 20 years clearly shows a decreasing tendency, as marked by the red dotted line.

Another result is, that ASes that are a hub today, have already been hubs for the last 20 years. We chose to plot historic ranks of the today top-5 ranked hubs (in terms of path-ratio) from 2000 until today in Figure 7.6.

Next to ranking ASes by their path-ratio we also investigate into essentiality of ASes. As outlined in Chapter 4 we assess the essentiality of Autonomous Systems by their essentiality-score. While we could already see, that hubs are not necessarily also ASes with a high essentiality-score, we present, how the essentiality-score of hubs changes over the years from 2000 until today. For this step we evaluate how many ASes per year depend on another AS with at least 5%, 10%, 20%, 30%, 40% and 50%. Consider an example: In one year, AS P has an essentiality-score of 20% for AS Q. AS R has an essentiality-score of 15% for AS S. Therefore, 2 ASes depend on another AS with at least 10%, 1 AS depends on another AS with at least 20%. As we want to compare this parameter over time, we take the relative amount of depending ASes. Consider that there are 100 ASes (in total) present on our example-dataset. Hence, 2% of ASes depend on another AS with at least 10%. 1% of ASes depend on another AS with at least 20%. The results for the last 20 years are shown in

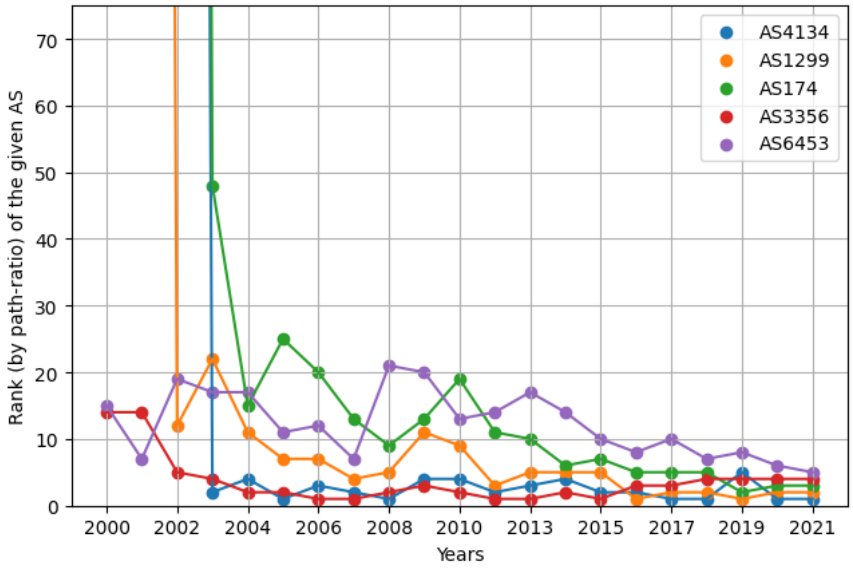

Figure 7.6.: It becomes clear, that ASes that hubs that are top-ranked today (in terms of path-ratio) have already been top-ranked in the past.

Figure 7.7. We can observe, that the percentage of ASes which are essential for other ASes shows a decreasing tendency over the years.

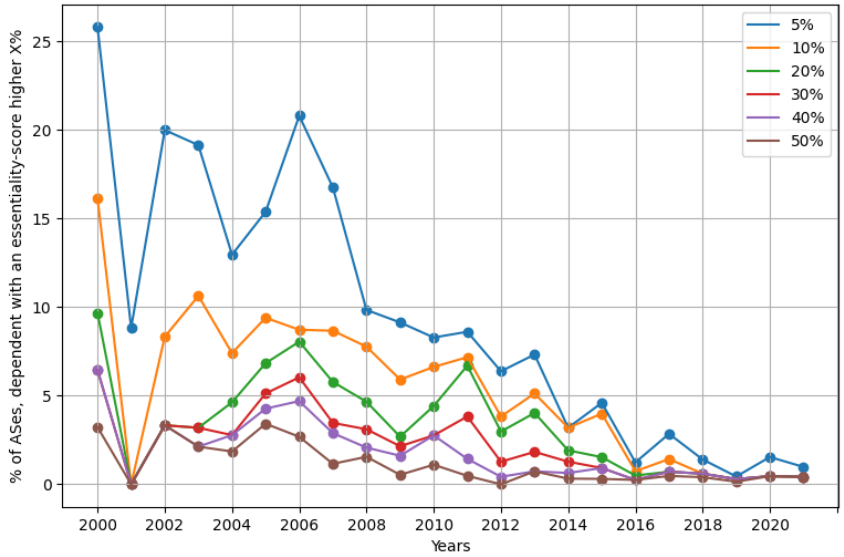

Figure 7.7.: Percentage of ASes that depend on other ASes with different essentiality-scores (5%, 10%, 20%, 30%, 40%, 50%)

Note, that results inferred from datasets having top-10 Tranco lists and having top-20 Tranco lists as path-origins go hand in hand. Therefore we only presented results gotten from datasets based on top-20 Tranco lists in this section.

7.2. Analyzing AS-paths from a local perspective of chosen countries

While maps, gotten for a global perspective of analysis show a high amount of paths, which result in a quite filled visualization, maps for a local perspective provide a much clearer view. Therefore it is possible to evaluate topology by looking at those maps. In the following, we provide maps for different countries on different continents. We proceed with describing observations we make in topology. Furthermore, we show, which countries act as important hosts for local routing of the given country. On top of that, we provide data that shows if there are essential ASes.
For all maps, markers are Autonomous Systems, edges represent AS-paths between ASes. Recall, that AS-paths are directed and go from path-origin towards path-destinations. Hence, data flows into the opposite direction. Markers with a red dot in their center represent path-origins, markers with a green dot in their center are path-destinations. The bigger a marker is, the higher is the path-ratio of the corresponding AS. Note, that all ASes that have a path-ratio higher than 5% have a label with their name on the map. Further information like path-ratio, additional locations and AS-names for ASes with a path-ratio lower than 5% can be viewed in the interactive version of our mapping-tool. As in previous sections we use different datasets for analysis of a local perspective, which are presented in the beginning of each subsection. Again, we always did analysis for two different datasets. As the resulting maps where almost equal, we only present maps gotten from datasets which are based on top-20 Tranco lists

7.2.1. Local perspective of Australia

The first country we perform an analysis for is Australia. Table 7.4 presents datasets for analysis of Australian topology and the resulting map for dataset 2021_oct_top20_aus is shown in Figure 7.8[1].
One observing that stands out is, that most path-origins are located in the United States and Asia. Another observing is, that hubs are not only located in Australia but also in

[1]Australian ISPs were found at ispquicklist.com [40] and resolved to the following ASNs: AS7575 AS55354 AS23552 AS24575 AS34393 AS15699 AS8893 AS9398 AS24215 AS4764 AS38611 AS24093 AS1221 AS137432 AS55581 AS136768 AS24033 AS7545 AS4817 AS38285 AS38263 AS4817 AS18398 AS17999 AS7545 AS31955 AS7579 AS45245 AS7546 AS37990 AS4764 AS10113 AS10115 AS18192 AS60258 AS10143 AS7718 AS135767 AS7604 AS4739 AS4802 AS16276 AS4739 AS38285 AS4826 AS137525 AS55834 AS45654 AS58538 AS4854 AS4802 AS7474 AS4643 AS4804 AS4740 AS37965 AS58857 AS10026 AS55988 AS17486 AS38263 AS38289 AS7477 AS51596 AS18390 AS37113 AS28787 AS28787 AS1221 AS132029 AS24130 AS4817 AS18398 AS7545 AS17999 AS38046 AS9543 AS50657 AS39093 AS196906 AS7604 AS18195

dataset	path origins	path destinations	collection date
2021_oct_top10_aus	Tranco top-10 pages	ASNs of Australian ISPs	2021-10-04
2021_oct_top20_aus	Tranco top-20 pages	ASNs of Australian ISPs	2021-10-04
...			
2000_oct_top10_aus	Tranco top-10 pages	ASNs of Australian ISPs	2000-10-04
2000_oct_top20_aus	Tranco top-20 pages	ASNs of Australian ISPs	2000-10-04

Table 7.4.: Datasets of AS-paths from 2000 to 2021 used in analysis with a local focus on Australian topology.

various other countries. Also the biggest markers and therefore the most important hubs are not located in Australia itself, but in Asia and Europe. Still, the highest amount of hubs locates to Australia. Next to Australia itself which hosts 39.5% of hubs, 30.2% of hubs are located in China and 11.6% of hubs locate to the United States. The rest of the hubs is located in various other countries. It also seems that (with an exception for Africa and South America) each continent has one major hub, that connects Australia to this continents countries. In Europe this is AS1299 which is located in Sweden. For Asia this hub is AS4134 which locates to China and for North America this is AS27281. Finally, Australia itself has multiple hubs which are shown in a closer perspective. (See Figure 7.9.) Those hubs are AS6453, AS4637, AS6939, AS7545, AS8075, AS6461 and AS174.

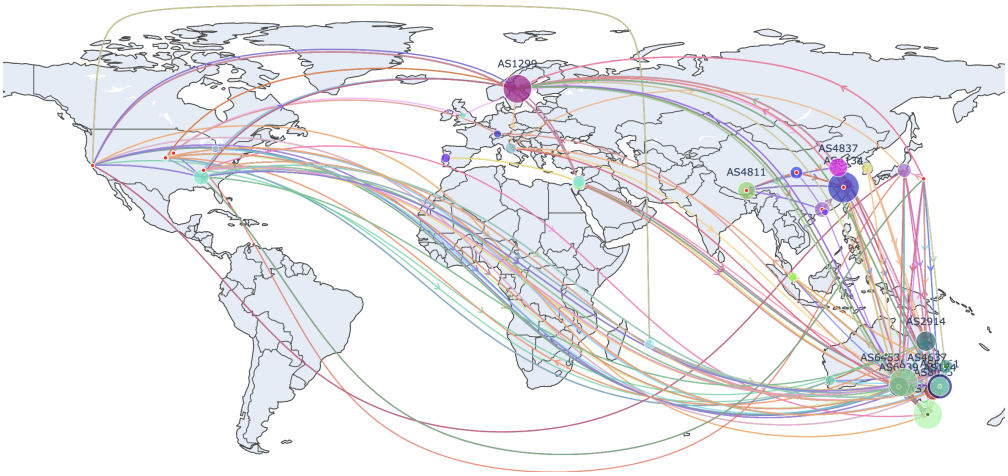

Figure 7.8.: A local perspective for Australia based on dataset 2021_oct_top20_aus. Markers are Autonomous Systems, edges are paths. Markers with a red dot in the center are path-origins, markers with a green dot in the center are path-destinations.

We also investigate if there are ASes from those hubs discovered in Figure 7.8 that are essential for routing. Figure 7.10 shows for how many from Australian ASes there is another Australian AS that has an essentiality-score of at least 5, 10, 20, 30, 40 and 50% for it. Data presented in this plot stems from every years top_20 dataset presented in Table 7.4. It turns out that we are not able to make out a clear tendency for essentiality of ASes over

the last years. What we are able to say is, that around 16% of Australians ASes depend on other Australian ASes with at least 10% of their paths, today. However there are no ASes which are dependent on other ASes with more than 20% of their paths.
Note, that we repeated analysis top_10 datasets, shown in Table 7.4 and got the same results as for top_20 datasets.

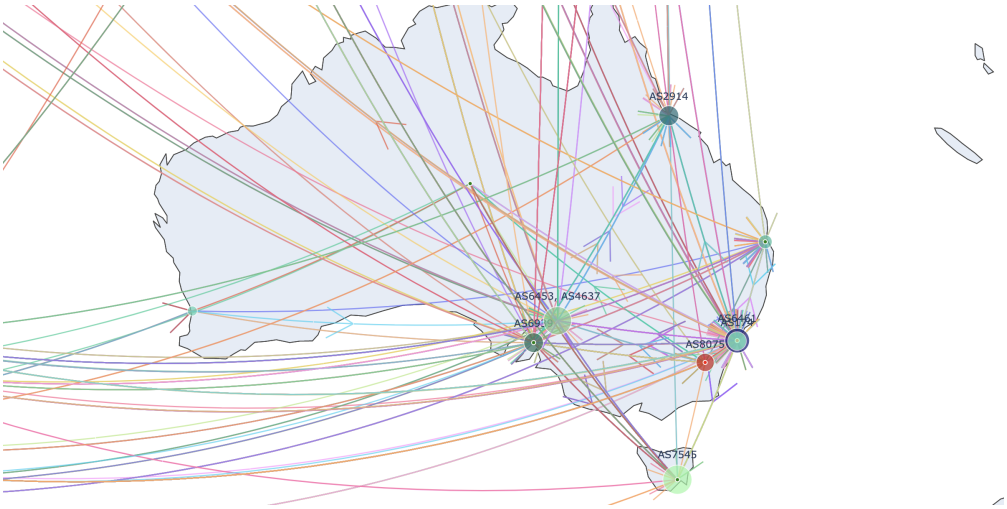

Figure 7.9.: A more detailed local perspective for Australia based on dataset 2021_oct_top20_aus

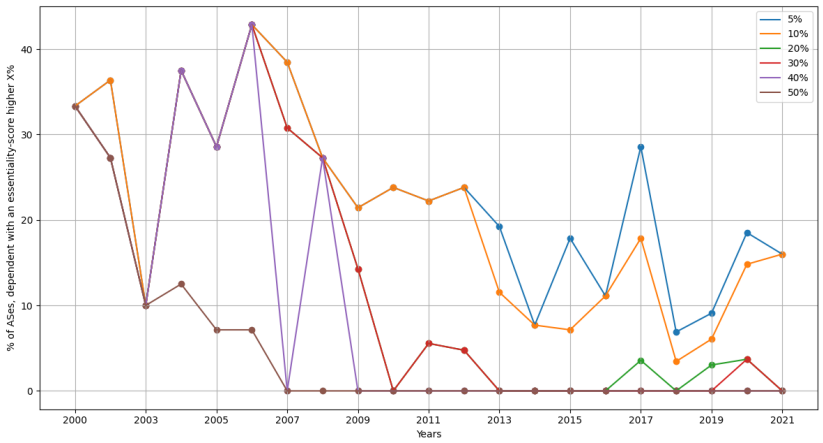

Figure 7.10.: Fraction of Australian ASes for which another Australian AS exists that has an essentiality-score of at least 5, 10, 20, 30, 40 and 50% for it.

dataset	path origins	path destinations	collection date
2021_oct_top10_usa	Tranco top-10 pages	ASNs of American ISPs	2021-10-04
2021_oct_top20_usa	Tranco top-20 pages	ASNs of American ISPs	2021-10-04
...			
2000_oct_top10_usa	Tranco top-10 pages	ASNs of American ISPs	2000-10-04
2000_oct_top20_usa	Tranco top-20 pages	ASNs of American ISPs	2000-10-04

Table 7.5.: Datasets of AS-paths from 2000 to 2021 used in analysis with a local focus on American topology

7.2.2. Local perspective of United States of America

Another country of interest for local analysis are the United States of America. Datasets processed for results are presented in Table 7.5[2] topology for the USA is shown in Figure 7.11 and builds on dataset 2021_oct_top20_usa.
One outstanding observation is, that hubs are located much more centralized to the United States than they were to Australia on the previous map. This seems to make sense, because there are no path-origins located in countries other than the USA and Asia. Following that fact, one would expect to see no paths to Europe. It surprises, that this is not the case. Viewing topology shows that even though there are no path-origins or path-destinations located in Europe there exist routes which originate in the USA take a hop over Europe and are routed back to the USA towards the path-destination. Such phenomena is called *routing-detour*. Considering the paths, it stands out, that there is a high amount of paths, which directly connects the United States with Asia and very little paths exist to the rest of the world. Next to the United States itself which hosts 55.8% of hubs, China hosts 25.6% of hubs. The remaining 18.6% of hubs are located in several other countries.
As for Australia, we also show essentiality of ASes for American topology. Figure 7.12 presents for how many American ASes there exist other American ASes that have essentiality-scores of at least 5, 10, 20, 30, 40 and 50% for them. We can identify a decreasing tendency for essential ASes over the last years. Despite this fact, there is a fraction of 5% of ASes which are dependent on other ASes with over 50% of their paths. An even higher fraction of ASes is dependent on other ASes with 10 to 20% of their paths.

7.2.3. Local perspective of China

China is the next country we run local analysis for. It seems to be an interesting candidate, because it is a country known for Internet-censorship and on the other hand we

[2] American ISPs were found at ispquicklist.com [40] and resolved to the following ASNs: AS7018 AS3561 AS20115 AS6181 AS16586 AS7922 AS22773 AS4355 AS13977 AS30036 AS18817 AS6128 AS6079 AS63365 AS7065 AS19108 AS32654 AS19406 AS46699 AS7221 AS32622 AS29930 AS21628 AS701 AS6167 AS2828 AS702 AS22394 AS7046 AS703 AS11486 AS12079 AS14551 AS15133 AS23148 AS6984 AS6066 AS704 AS705 AS6256 AS14210 AS12234 AS16724 AS12083 AS7155 AS16491 AS7168 AS7029 AS6983 AS7062 AS11776 AS23100 AS36790 AS12035 AS21688 AS32808

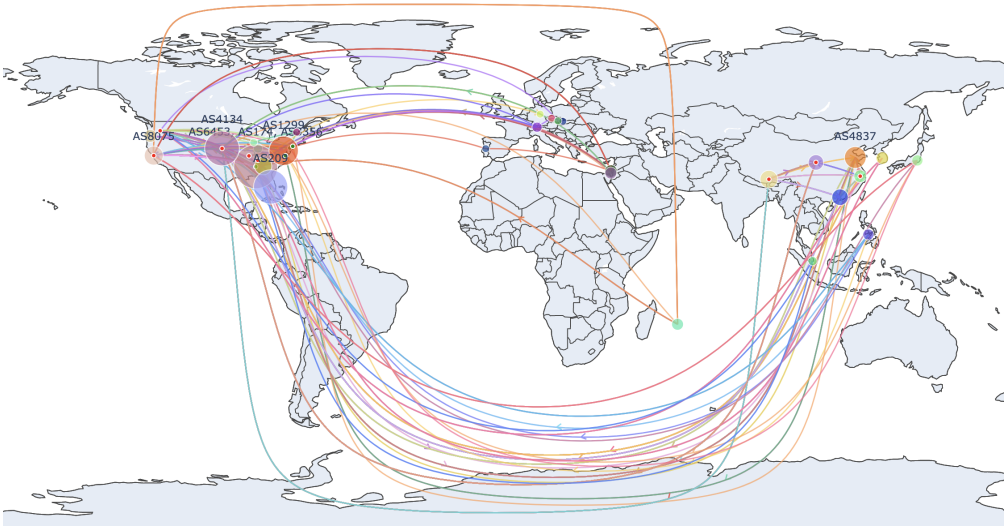

Figure 7.11.: A local perspective for the United States of America based on dataset 2021_oct_top20_usa. Markers are Autonomous Systems, edges are paths. Markers with a red dot in the center are path-origins, markers with a green dot in the center are path-destinations.

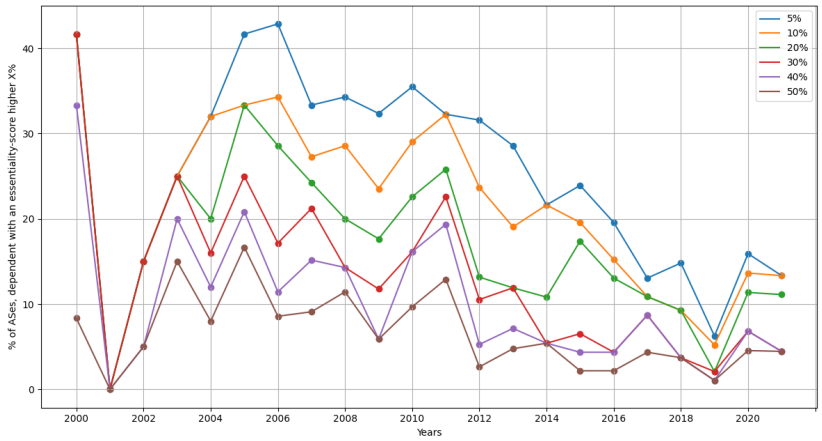

Figure 7.12.: Fraction of American ASes for which other American ASes exist that have an essentiality-score of at least 5, 10, 20, 30, 40 and 50% for them.

saw it playing a central role for topology of countries, analyzed in previous steps. A map for the Chinese perspective is provided in Figure 7.13 and Table 7.6 shows datasets used for analyzing Chinese topology[3]. As already observed for previous countries, most path-origins are located in the United States and China. This also causes the fact, that there is a

[3]Chinese ISPs were found at ispquicklist.com [40] and resolved to the following ASNs: AS4808 AS4847 AS4837 AS4814 AS17816 AS58453 AS56042 AS9808 AS56040 AS58453 AS140895 AS4134 AS4809 AS4812 AS4809 AS23764 AS9394 AS45069 AS63711

dataset	path origins	path destinations	collection date
2021_oct_top10_china	Tranco top-10 pages	ASNs of Chinese ISPs	2021-10-04
2021_oct_top20_china	Tranco top-20 pages	ASNs of Chinese ISPs	2021-10-04
...			
2000_oct_top10_china	Tranco top-10 pages	ASNs of Chinese ISPs	2000-10-04
2000_oct_top20_china	Tranco top-20 pages	ASNs of Chinese ISPs	2000-10-04

Table 7.6.: Datasets of AS-paths from 2000 to 2021 used in analysis with a local focus on Chinese topology

high amount of paths, which stays in China and directly connects path-origins and path-destinations. One interesting detail is, that even though only one path-origin is located in Europe, with AS1299 there is a major hub on this continent. The same goes for AS37662 which is located in Somalia despite the fact that there is no path-origin located in Africa. Taking paths into account which are routed via those hubs it becomes clear, that those hubs connect Asia with North America. Next to paths, which take those 'stopovers' there is a high amount of paths, directly connecting Asia and North-America. Most of the local hubs (42.5%) are located in China. Additional 20% of hubs are located in the United States. An interesting fact is, that 15% of hubs for Chinese routing are located in Australia. This stands out, because there are no path-origins located in Australia. The remaining hubs are scattered over various other countries. As for other countries we also investigate into essentiality of Chinese ASes. Again, we plot for how many of Chinas ASes there are other Chinese ASes which have an essentiality-score higher than 5, 10, 30, 50, 80 and 90% for them. Note, that in contrast to other countries we raised the level of plotted essentiality-scores up to 90%, as we could observe that there is a small fraction of approximately 3% of ASes which depend on other ASes with more than 90% of their paths. What also stands out is, that until 2011 we are unable to detect any essential ASes in Chinese topology. These results are shown in Figure 7.14.

7.2.4. Local perspective of Germany

Next to aforementioned countries we also analyze topology of our home country Germany. Datasets used for analysis of German topology are presented in Table 7.7[4]. The

[4]German ISPs were found at ispquicklist.com [40] and resolved to the following ASNs: AS8881 AS211357 AS60294 AS15763 AS12312 AS9145 AS35265 AS39356 AS13045 AS8422 AS16097 AS20825 AS3209 AS15943 AS3209 AS8426 AS3320 AS5430 AS8767 AS6805 AS47309 AS6724 AS16202 AS42652 AS3320

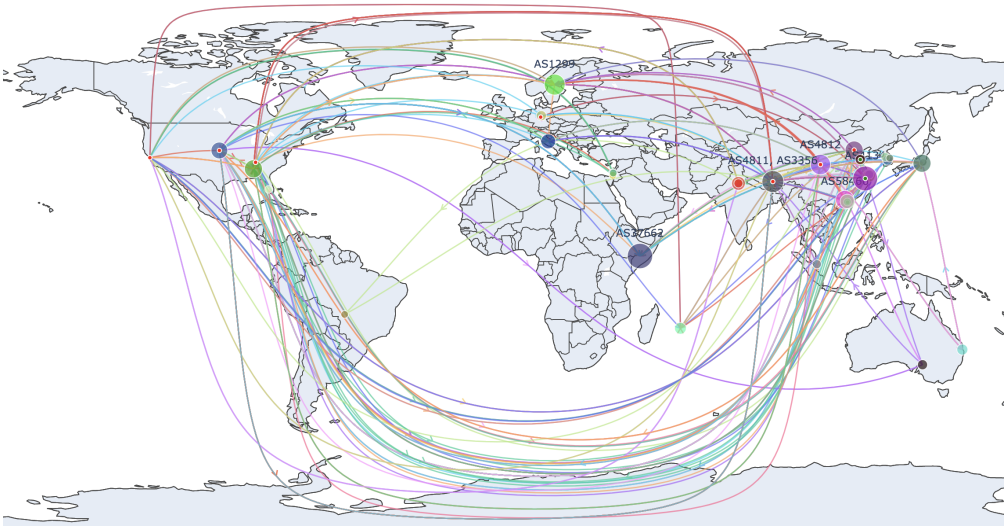

Figure 7.13.: A local perspective for China based on dataset 2021_oct_top20_china. Markers are Autonomous Systems, edges are paths. Markers with a red dot in the center are path-origins, markers with a green dot in the center are path-destinations.

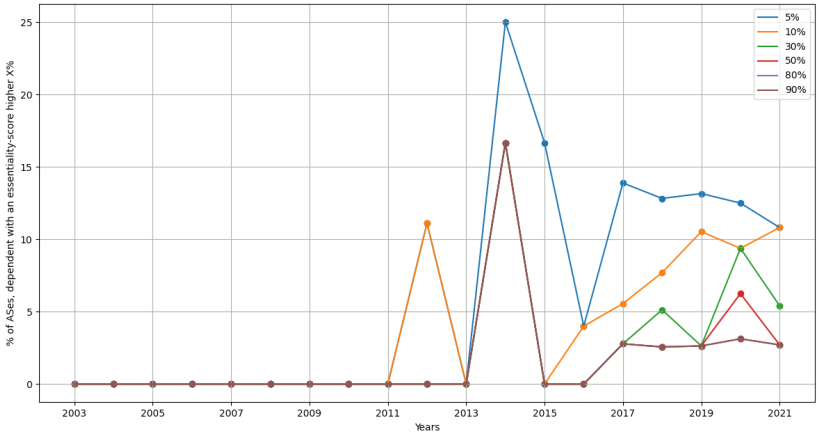

Figure 7.14.: Fraction of Chinese ASes for which other Chinese ASes exist that have an essentiality-score of at least 5, 10, 30, 50, 80 and 90% for them.

resulting map, created in the process of analysis is based on dataset 2021_oct_top20_ger. See Figure 7.15 for the result.

One can observe, that most of the hubs are also located in Germany. Next to this, we also find AS4134 with a path-ratio of 26.5% the biggest hub, located in China. China also is the country, which hosts with 27.9% of hubs most of the hubs for German topology, next to Germany itself (also 27.9%). Next to those two, the United States host 11.6% of hubs followed by Australia with 7%. An interesting result that stands out is, that again routing-detours can be observed when taking a closer look at Germany. (See Figure 7.16.)

dataset	path origins	path destinations	collection date
2021_oct_top10_ger	Tranco top-10 pages	ASNs of German ISPs	2021-10-04
2021_oct_top20_ger	Tranco top-20 pages	ASNs of German ISPs	2021-10-04
...			
2000_oct_top10_ger	Tranco top-10 pages	ASNs of German ISPs	2000-10-04
2000_oct_top20_ger	Tranco top-20 pages	ASNs of German ISPs	2000-10-04

Table 7.7.: Datasets of AS-paths from 2000 to 2021 used in analysis with a local focus on German topology

AS8075, a path-origin clearly has outgoing AS-paths which are not directly leading towards one of the German path-destinations. Hence, even though both, path-origin and path-destinations are located in Germany they are not directly connected via a path of one hop, as one would expect. Instead, paths originated in AS8075 leave Germany, take one or more hops via an AS located in another country and are routed back to Germany. Also

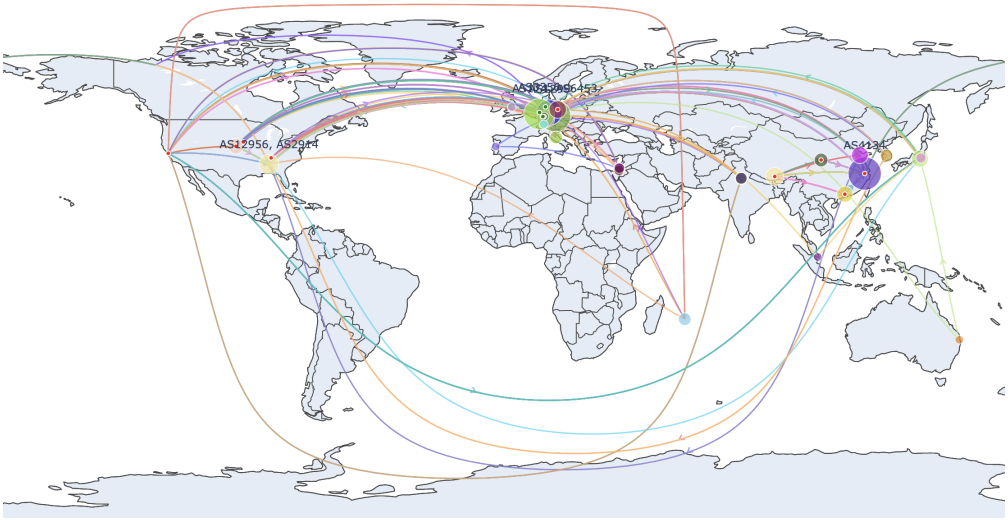

Figure 7.15.: A local perspective for Germany based on dataset 2021_oct_top20_ger. Markers are Autonomous Systems, edges are paths. Markers with a red dot in the center are path-origins, markers with a green dot in the center are path-destinations.

for German topology we present results on essential ASes. Again, we plot for how many German ASes there are other German ASes that have essentiality scores of at least 5, 10, 30, 50, 80 and 90% for them. Data can be found in Figure 7.17. There's one point in this data which surprises, in view of the fact that Germany is not a country with Internet censorship: Over the time there are some ASes which depend on other ASes with more than 90% of their paths. However the number of those highly dependent ASes seems to decrease over the last 8 years. Despite that development there still is a fraction of approximately 17% of German ASes which depends on other ASes with more than 5% of their paths.

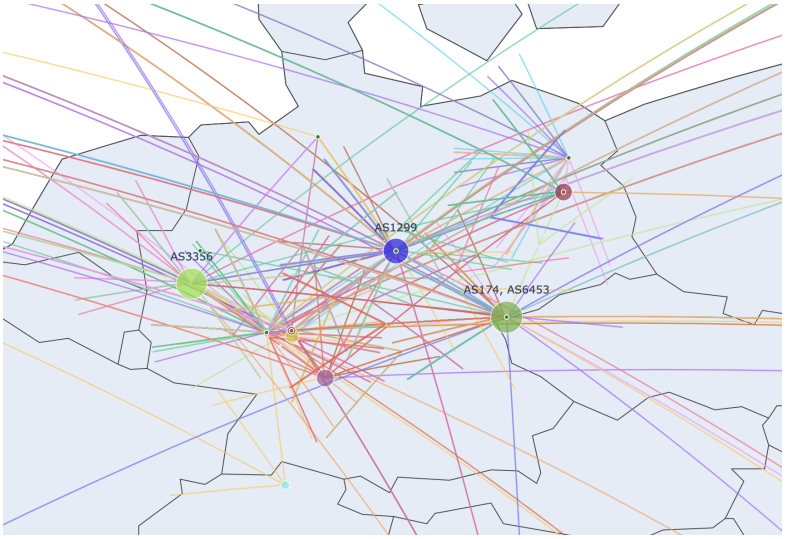

Figure 7.16.: A closer perspective of German topology based on dataset 2021_oct_top20_ger. AS8075, a path-origin clearly has routes which take detours to other countries, even though destinations are located in Germany, too.

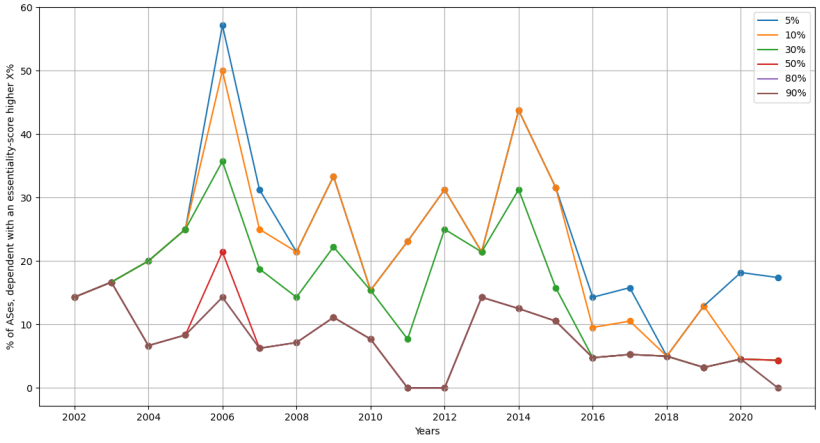

Figure 7.17.: Fraction of German ASes for which other German ASes exist that have an essentiality-score of at least 5, 10, 30, 50, 80 and 90% for them.

7.3. Analyzing misconfigurations in historical BGP-data

As already outlined in Section 2.5, there have been events in history where misconfigurations lead to rerouting and worldwide unreachability of certain services. Based on historical AS-paths we analyzed topology for two of those events.

7.3.1. Misconfiguration of Pakistan Telecom 2008

On February 24, 2008 Pakistan Telecom (AS17557) started to announce one of YouTube's (AS36561) IP-prefixes. As they announced a more specific prefix than YouTube itself, Autonomous Systems worldwide adopted the route announced by Pakistan Telecom. Ripe gives a detailed report on this incident on its blog [41]. They also provide a timeline of YouTube's reactions and how they tried to regain control over their prefix. In the following we give a brief overview on this incident. The most relevant points are shown in Figure 7.18. We numbered the time-ranges between those points and provide maps for each of those time-ranges to observe consequences of the shown actions. Red AS-paths lead to YouTube's AS (AS36561), while green AS-paths lead to Pakistan Telecom's AS (AS17557). The following events mark important points of the incident:

1. Before 18:47 YouTube's AS (AS36561) announces YouTube's IP-prefix `208.65.152.0/22`. (See Figure 7.19.)

2. At 18:47 Pakistan Telecom's AS (AS17557) starts to announce the prefix `208.65.153.0/24`. As this prefix is a more specific prefix of a portion of YouTube's previously announced prefix, ASes worldwide start to drop routes to YouTube's originally announced prefix. Instead, they route their traffic towards Pakistan Telecom. However, some valid AS-paths to YouTube still remain. One reason might be, that these ASes had static AS-paths applied. This means that they configured a direct connection to YouTube's AS in their router and did not get the wrong paths, provided via BGP. (See Figure 7.20.)

3. YouTube learns of this incident and tries to regain control over its prefix by also announcing `208.65.153.0/24`. This step does not show the desired effect. So at 20:18 YouTube announces two more specific prefixes than Pakistan Telecom, `208.65.153.0/25` and `208.65.153.128/25`, which together are equivalent to `208.65.153.0/24`. This step brings YouTube back in control of some AS-paths towards the hijacked prefix. Still there are some paths which lead to Pakistan Telecoms's AS (See Figure 7.21.)

4. At 21:01, Pakistan Telecom's provider stops to pass the BGP-announcement towards the hijacked prefix on and therefore the wrongly announced AS-paths disappear. (See Figure 7.22)

Since historical data is rare, there are only few paths observable, but nonetheless one can clearly monitor the incident.

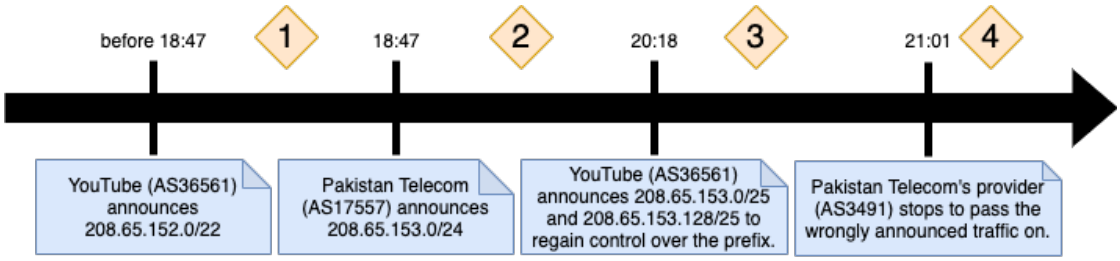

Figure 7.18.: The most relevant points for routing during the Pakistani hijack of YouTube's IP-prefix on February 24, 2008.

Figure 7.19.: Time-range 1: State before Pakistan's misconfiguration. The large black dot is YouTube's AS36561 which originates all shown AS-paths. All other dots represent ASes that are destinations of YouTube's routes.

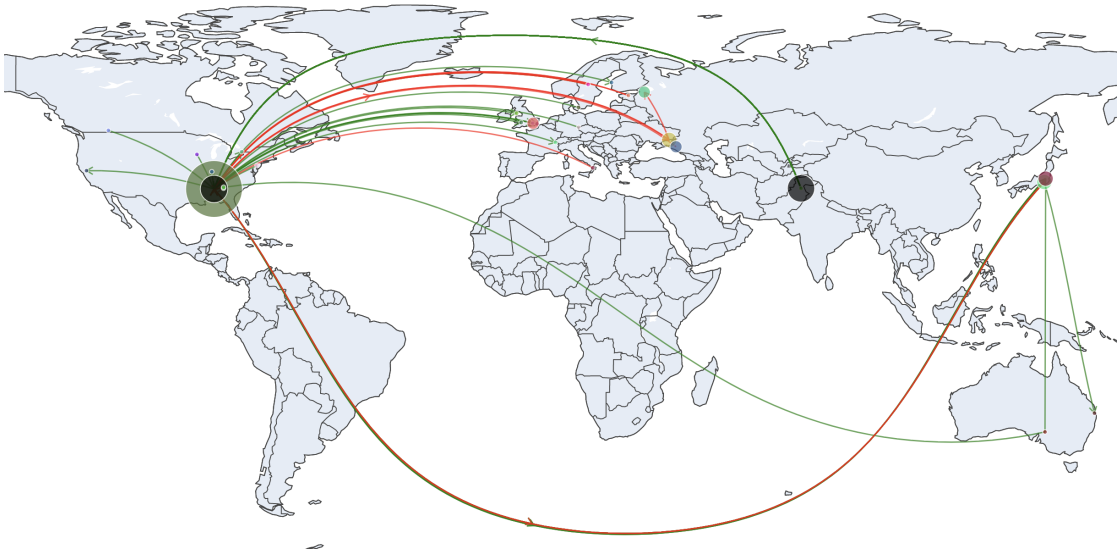

Figure 7.20.: Time-range 2: State after Pakistan Telecom announces YouTube's prefix. The additional black dot is Pakistan Telecom's AS17557, which announces YouTube's prefix. As one can see, a high fraction of YouTube's paths (red) is taken over by Pakistan Telecom's paths (green).

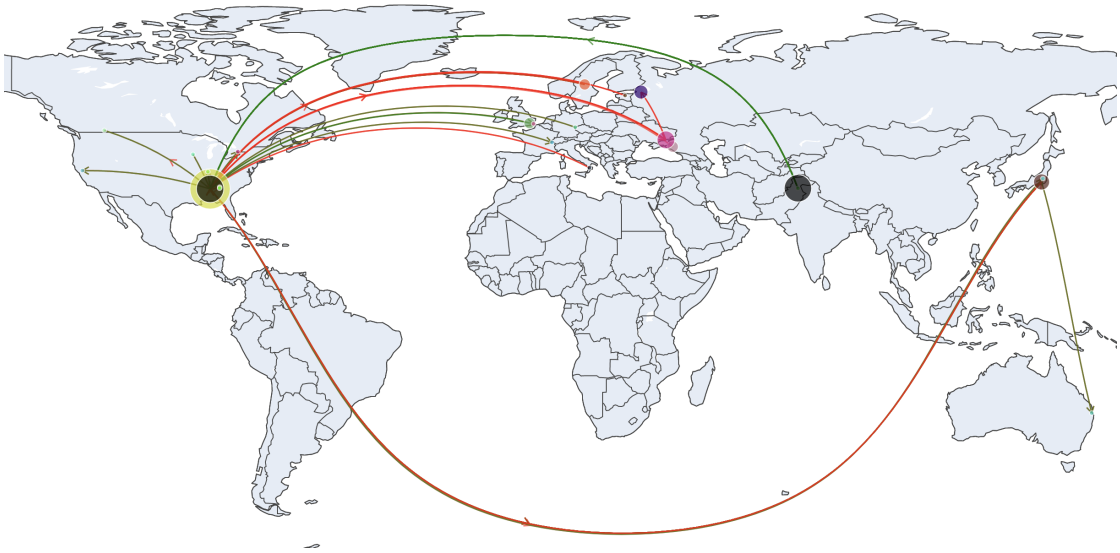

Figure 7.21.: Time-range 3: YouTube announces two more specific prefixes than Pakistan Telecom. Pakistan Telecom's paths are still active, but YouTube's paths are adopted by other ASes again.

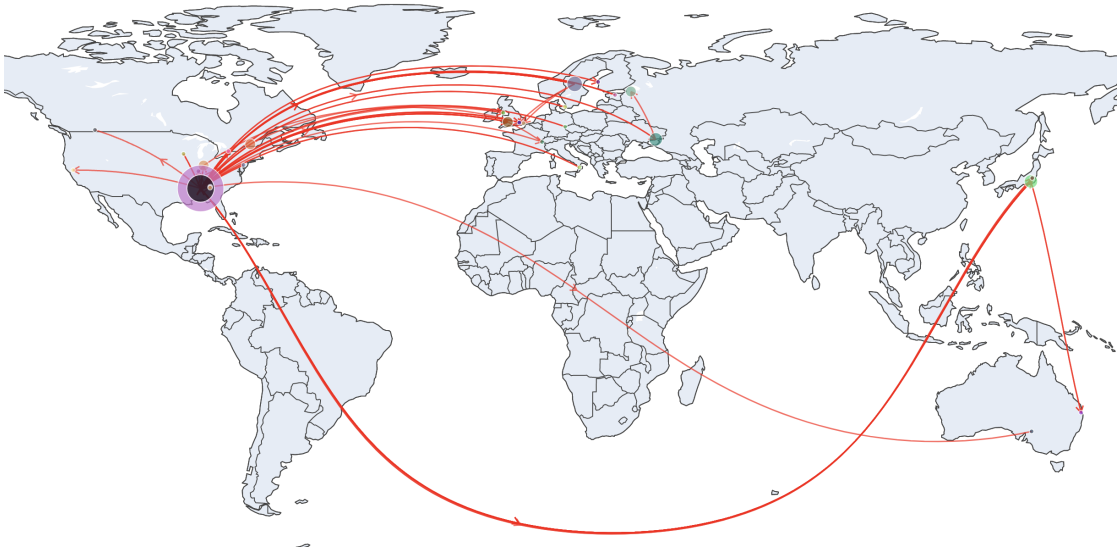

Figure 7.22.: Time-range 4: Pakistan Telecom's provider stops to pass on the wrongly announced AS-paths. Their routes are gone and routing is back at its initial state.

7.3.2. Facebook Outage 2021

On October 4, 2021 Facebook and its various Services (WhatsApp, Instagram, ...) suffered a worldwide outage, lasting around five hours. As Facebook reports on their official blog, this outage was caused by an internal misconfiguration [42]. This misconfiguration caused their infrastructure to drop almost all their AS-paths to their DNS servers. One of Facebook's affected subnetworks that was (almost) not reachable during this incident anymore, is `185.89.218.0/23`. Interesting points of time during this incident are provided by Cloudfare [13]. We created a timeline based on Cloudfare's observations (see Figure 7.23) and mapped topology at the most interesting points during the incident. We numbered the time-ranges between those points and provide maps for each of those time-ranges to observe consequences of the shown actions. The following events mark important points of the outage:

1. Before 15:40 Facebook's AS (AS32934) announces their IP-prefix `85.89.218.0/23`. (See Figure 7.24.)
2. At 15:40 Facebook's AS (AS32934) drops almost all paths to several of their subnetworks, including `85.89.218.0/23`. (See Figure 7.25.)
3. At 21:20, after several hours, Facebook's starts to announce the paths to `85.89.218.0/23` again. (See Figure 7.26)

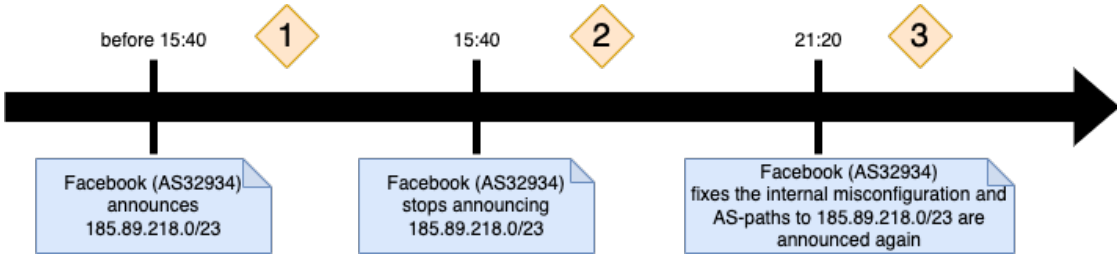

Figure 7.23.: The most relevant points during the Facebook outage on October 4, 2021.

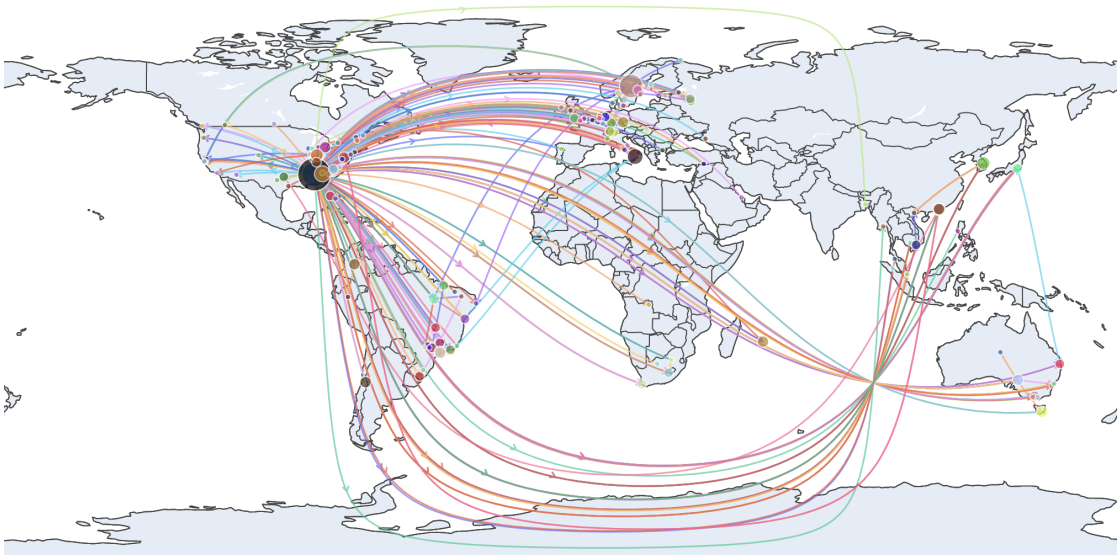

Figure 7.24.: State before Facebook's outage. The large black dot is Facebook's AS32934 which originates all shown routes. All other dots represent ASes that are destinations of Facebook's routes.

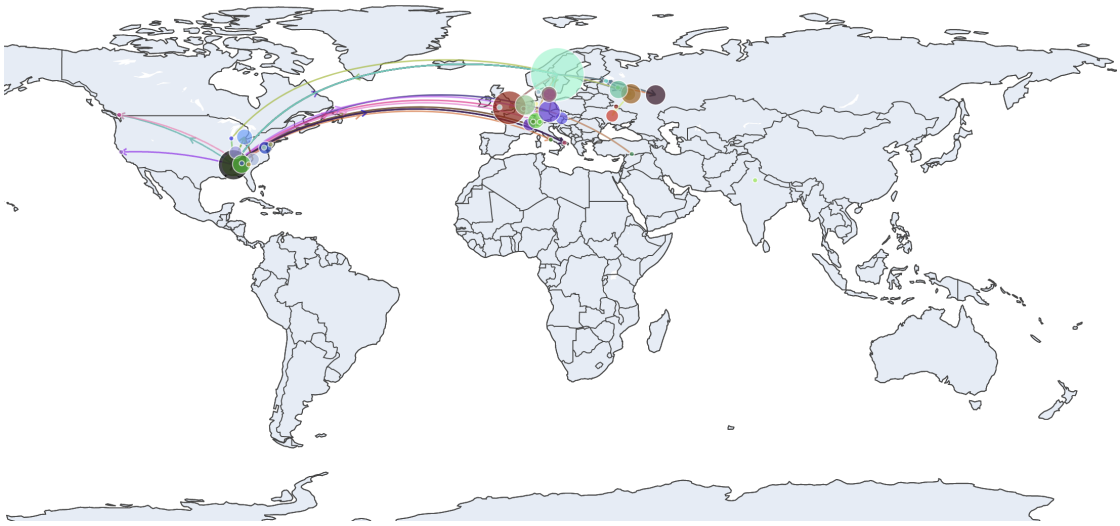

Figure 7.25.: State during Facebook's outage. One can observe that there are only few AS-paths are present, because the high majority was dropped by their own AS (AS32934).

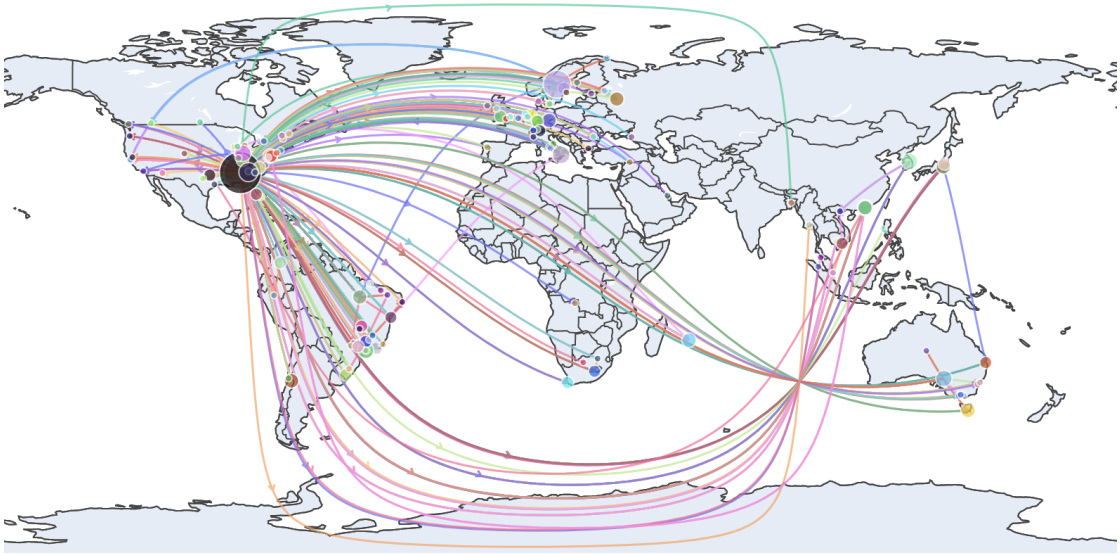

Figure 7.26.: State after the Facebook's outage. AS-paths are announced by AS32934 as they were before the incident. Facebook's routing information are back at their initial state and their services can be reached again.

8. Discussion

This section discusses results from Chapter 7 and aims to identify tendencies that can be inferred from them.

8.1. Global hubs

There are some highly interesting tendencies that can be derived from presented data. Given the fact, that some ASes are present on more than 10% of our datasets AS-paths and one of them even intercepts 25% of paths is one indicator that there are indeed some highly important ASes which act as Hubs of Control for global routing. Figure 7.1 showed, that Hubs of Control are mainly located in the United States and China. This observation matches results from earlier research that also identified the USA and China as essential countries for global routing [25]. As China is a country, known for Internet-censorship, this fact could raise concerns. One surprising fact is, that after the USA and China, Switzerland is the third ranked country in terms of number of hubs. This result stems from the fact, that we counted the amount of hubs, present among countries, but did not take into account how big those hubs are. Hence, a country with a high amount of smaller hubs, will strike out countries which have a single major hub. Despite this inaccuracy we argue, that the results presented in Figure 7.1 provide a decent measure for which countries got a high amount of power in global topology.

One result that raises attention is, that the amount of hubs stayed the same over the last years. (See Figure 7.4.) On top of that, the worldwide total amount of ASes showed a continuing growth over the past years. This growth does not seem to come to an end in the near future. Consequently, the power of the few existing hubs already increased over the last years. We already showed this tendency in Figure 7.5. We also showed, that todays top-ranked ASes have already been top-ranked for the last 10 to 20 years. Thus, if the outlined tendencies continue over the next years, todays hubs will increase their power on global routing.

One finding that facilitates is, that we were not able to make out a set of Autonomous Systems that is essential for routing. We could indeed observe, that there appear some ASes that are highly dependent on other ASes every once in a while. These observations did not resume over the years. We suppose that in those cases newly registered ASes, which had only few providers in their early days appeared in our dataset. Over the time those ASes expanded their set of providers and essentiality decreased. Consequently, we can not find a tendency which accounts that very powerful ASes could not be circumvented. This result also matches the insights of Edmundson et al., who describe, that in nearly all cases it is possible to route around certain locations [27]. Another interesting insight is the development of numbers of ASes which are dependent on other ASes. We showed this trend in

Figure 7.7. In the early 2000s there was indeed a high fraction (up to 25%) of ASes in high dependence of other ASes, but we observed a decreasing tendency over the last 20 years. We argue that this trend comes from the growing number of ASes around the globe. The more Autonomous Systems exist around the world, the more routes emerge. Hence, there also appear more alternative routes for a given set of path-origins and path-destinations. Consequently the dependency of ASes on certain other ASes decreases.

8.2. Local hubs

In Section 7.2 we ran analysis for four different countries. One interesting result is, that most path-origins (recall that path-origins are the ASes of top-ranked websites) are located in the United States and Asia. We already observed, that despite of this fact, routes from and to the United States and China still take hops over countries which do not host path-origins or path-destinations. We suppose that this behavior is caused by business-relationships between ASes. E.g. a path-origin and a path-destination both located in the United State do not have a common peer-to-peer-relationship. The origin peers with an AS in Europe which also exchanges traffic with the destination-AS in the United States for free. Therefore the AS-path does not directly take a hop from origin to destination, but via the AS located in Europe.
We could observe, that hubs for local routing are similar to ASes which are hubs for global routing. This fact holds for all of the four countries we analyzed. Again, we noticed that the United States and China are the locations where most of those hubs are located. Hence, also for local routing the United States and China are heavily important countries. One could question, why Germany which hosts the DE-CIX, the worlds biggest IXP (in terms of average traffic rate [8]), does not play a prominent role in routing. This is, because at an IXP ASes are connected with each other. Hence, it does not mean, that ASes must be located directly at the IXPs location.
Essentiality-scores of American and Chinese topology show, that there are some ASes (around 5%) which are completely dependent on other ASes. It surprises that there exist ASes in the United States which are completely dependent on other ASes, as the USA are not a country where Internet-censorship is deployed on BGP-level. In contrast to that, we would expect far more ASes being far more dependent in Chinese topology, as China is a country, actively deploying Internet-censorship on BGP-level. Next to China and the USA we could not observe ASes for Australian and German topology which are entirely dependent on other ASes. In Australian topology this fact also applies for the last 10 years whereas in Germany the number of entirely dependent ASes decreased to zero over the last years. We investigated into the affected German ASes and found out, that those dependent ASes belonged to the same organization. Hence, it is questionable, if we can speak of dependencies, if ASes of the same owner are dependent on each other. However we infer, that for all of the analyzed countries there is no high enough fraction of ASes, completely dependent on other ASes that we could speak of essential ASes.

8.3. Insights from global and local hubs

Considering our previous discussions of a local and a global perspective of hubs, it becomes clear that insights and tendencies are similar for local and global perspectives. It is well known, that the United States and China are great political rivals. In view of this fact it is especially interesting, that those two countries are the major global players in the Internet. Taking Americas recent past regarding online-privacy, surveillance of mobile devices and violations of digital rights into account, the fact that the USA of all countries is the major player in the Internet, leaves a bitter taste. What reliefs is, that of all hubs that are located in those countries, none are essential. Yet, there is no clear tendency concerning essential ASes for the future to detect. Looking ahead monitoring essentiality of ASes should become one key suspect of interest.
One question that remains is, if our dataset on which the results are based can be considered a representative fraction of the Internet. In Chapter 6 we outlined, why our dataset is limited. Compared to the work of Acharya et al. who chose to hypothesize on paths, our results confirm their findings. [4] This and the fact that their work is based on a dataset which seems more complete than ours seems to be an indicator, that findings from our dataset are equal to the ones from a complete dataset. However, it is not sure if their dataset represents AS-topology accurately as they hypothesize on paths. As both approaches bear their flaws, the finding of a more complete and also accurate dataset should be part of future work.

8.4. Misconfigurations

Next to analysis of global and local hubs, we were also able to visualize two incidents of rerouting caused by misconfigurations. The first incident presented in Section 7.3.1 is only one of several similar events where a misconfiguration happened directly on AS-level.
The second incident we analyzed, was not caused by an immediate misconfiguration on AS-level, but consequences were observable in BGP data. By analyzing this incident we underlined the viability of our methodology, once more.
Whether AS-paths are wrongly announced on purpose or not, it is questionable why it is still possible that this behavior occurs. There already exist attempts towards a solution of this problem (like BGPsec or RPKI), but it is questionable whether these approaches will be adopted by the majority of ASes around the globe.
Next to monitoring of essentiality of ASes the focus on securing the authenticity of AS-paths should be focused on in the future.

9. Conclusion and Outlook

In this thesis we gave an introduction to the essential protocol that connects Autonomous Systems around the globe: The Border Gateway Protocol. Furthermore we outlined, what could possibly go wrong by presenting different attack types and misconfigurations that may happen by mistake. Based on existing research and resources we worked out a methodology which investigates into global and local hubs. We used our methodology to run a detailed analysis with global and local focuses. Next to this analysis we were able to conclude on essential ASes. To the best of our knowledge we are the first ones who not only considered current data for analysis but also used historical BGP-data to infer trends in analysis. To demonstrate further usage of our tool we closed our findings by visualizing two severe misconfigurations, which lead to major traffic disruptions in the global Internet.

Applying our methodology on existing BGP-data showed, that we had only access to a fraction of the worlds AS-paths. Future work should deal with the challenge of collecting more AS-paths. Next to this, we suppose to keep an eye on the development of hubs and especially essential ASes. While we were not able to find ASes which could be considered 'essential' this might change in the future and if it does this should raise attention.

Bibliography

[1] A. Shah, R. Fontugne, and C. Papadopoulos, "Towards characterizing international routing detours," in *Proceedings of the 12th Asian Internet Engineering Conference*, ser. AINTEC '16, Bangkok, Thailand: Association for Computing Machinery, 2016, pp. 17–24, ISBN: 9781450345521. DOI: 10.1145/3012695.3012698. [Online]. Available: https://doi.org/10.1145/3012695.3012698.

[2] A. Shah and C. Papadopoulos, "Characterizing international bgp detours," 2015.

[3] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "RAPTOR: Routing attacks on privacy in tor," in *24th USENIX Security Symposium (USENIX Security 15)*, Washington, D.C.: USENIX Association, Aug. 2015, pp. 271–286, ISBN: 978-1-939133-11-3. [Online]. Available: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/sun.

[4] H. B. Acharya, S. Chakravarty, and D. Gosain, "Few throats to choke: On the current structure of the internet," *CoRR*, vol. abs/1806.07038, 2018. arXiv: 1806.07038. [Online]. Available: http://arxiv.org/abs/1806.07038.

[5] *World - autonomous system number statistics*. [Online]. Available: https://www-public.imtbs-tsp.eu/~maigron/RIR_Stats/RIR_Delegations/World/ASN-ByNb.html.

[6] Lixin Gao and J. Rexford, "Stable internet routing without global coordination," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, 2001.

[7] A. S. Tanenbaum, *Computer networks*. Upper Saddle River, NJ: Prentice Hall, 1996, p. 355.

[8] P. C. H. (PCH), *Internet exchange directory*, 2022. [Online]. Available: https://www.pch.net/ixp/dir#!mt-sort=avg%2Cdesc!mt-pivot=avg.

[9] R. Bush, R. Volk, and J. Heitz, "Bgp rpki-based origin validation on export," 2020.

[10] M. Lepinski and K. Sriram, "Bgpsec protocol specification," *RFC*, vol. 8205, pp. 1–45, 2017.

[11] D. McCullagh, "How pakistan knocked youtube offline (and how to make sure it never happens again)," *cnet.com*, 2008. [Online]. Available: https://www.cnet.com/news/how-pakistan-knocked-youtube-offline-and-how-to-make-sure-it-never-happens-again/.

[12] D. G. Blogs, "Internet-wide catastrophe - last year," *blogs.oracle.com*, 2005. [Online]. Available: https://blogs.oracle.com/internetintelligence/internet-wide-catastrophe%e2%80%94last-year.

[13] T. S. Celso Martinho, "Https://blog.cloudflare.com/october-2021-facebook-outage/," *The Cloudfare Blog*, 2021. [Online]. Available: https://blog.cloudflare.com/october-2021-facebook-outage/.

[14] C. D. Marsan, "Six worst internet routing attacks, How youtube, yahoo and others fell prey to router incidents and accidents," *networkworld.com*, 2009. [Online]. Available: https://www.networkworld.com/article/2272520/six-worst-internet-routing-attacks.html.

[15] B. Donnet and T. Friedman, "Internet topology discovery: A survey," *IEEE Communications Surveys & Tutorials*, vol. 9, pp. 56–69, 2007.

[16] P. Winter, R. Padmanabhan, A. King, and A. Dainotti, "Geo-locating bgp prefixes," *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 9–16, 2019.

[17] L. Gao and J. Rexford, "Stable internet routing without global coordination," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, 2001. DOI: 10.1109/90.974523.

[18] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, 2001. DOI: 10.1109/90.974527.

[19] *Route views*. [Online]. Available: http://routeviews.org.

[20] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: From art to science," in *2009 Cybersecurity Applications Technology Conference for Homeland Security*, 2009, pp. 205–211. DOI: 10.1109/CATCH.2009.38.

[21] *Archipelago*. [Online]. Available: https://caida.org/projects/ark.

[22] R. Winter, "Modeling the internet routing topology - in less than 24h," in *Proceedings of the 2009 ACM/IEEE/SCS 23rd Workshop on Principles of Advanced and Distributed Simulation*, ser. PADS '09, USA: IEEE Computer Society, 2009, pp. 72–79, ISBN: 9780769537139. DOI: 10.1109/PADS.2009.17. [Online]. Available: https://doi.org/10.1109/PADS.2009.17.

[23] *Ucla internet research lab*. [Online]. Available: https://irl.cs.ucla.edu.

[24] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson, "Internet atlas: A geographic database of the internet," in *HotPlanet '13*, 2013.

[25] J. Karlin, S. Forrest, and J. Rexford, "Nation-state routing: Censorship, wiretapping, and bgp," *ArXiv*, vol. abs/0903.3218, 2009.

[26] A. Edmundson, R. Ensafi, N. Feamster, and J. Rexford, "Nation-state hegemony in internet routing," *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 2018.

[27] ——, "Ran : Routing around nation-states," 2017.

[28] H. Roberts, D. Larochelle, R. Faris, and J. Palfrey, "Mapping local internet control," 2011.

[29] M. Wählisch, T. C. Schmidt, M. de Brün, and T. Häberlen, "Exposing a nation-centric view on the german internet – a change in perspective on as-level," in *Passive and Active Measurement*, N. Taft and F. Ricciato, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 200–210, ISBN: 978-3-642-28537-0.

[30] D. Oh and K.-W. Lee, "Study on the characteristics of the korea internet as-level topology using node degree and node connectivity metrics," *The Journal of Korean Institute of Communications and Information Sciences*, vol. 38, pp. 417–426, 2013.

[31] S. Zhou and G.-Q. Zhang, "Chinese internet as-level topology," *IET Commun.*, vol. 1, pp. 209–214, 2007.

[32] R. Fanou, P. Francois, and E. Aben, "On the diversity of interdomain routing in africa," in *Passive and Active Measurement*, J. Mirkovic and Y. Liu, Eds., Cham: Springer International Publishing, 2015, pp. 41–54, ISBN: 978-3-319-15509-8.

[33] O. Nordström and C. Dovrolis, "Beware of bgp attacks," *Comput. Commun. Rev.*, vol. 34, pp. 1–8, 2004.

[34] H. Birge-Lee, L. Wang, J. Rexford, and P. Mittal, "Sico: Surgical interception attacks by manipulating bgp communities," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19, London, United Kingdom: Association for Computing Machinery, 2019, pp. 431–448, ISBN: 9781450367479. DOI: 10.1145/3319535.3363197. [Online]. Available: https://doi.org/10.1145/3319535.3363197.

[35] V. L. Pochat, T. V. Goethem, and W. Joosen, "Rigging research results by manipulating top websites rankings," *CoRR*, vol. abs/1806.01156, 2018. arXiv: 1806.01156. [Online]. Available: http://arxiv.org/abs/1806.01156.

[36] V. L. Pochat, T. V. Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, *Tranco: A research-oriented top sites ranking hardened against manipulation.* [Online]. Available: tranco-list.eu.

[37] *Maxmind geolite2 free geolocation data.* [Online]. Available: https://dev.maxmind.com/geoip/geolite2-free-geolocation-data.

[38] Plotly, *Plotly python open source graphing library.* [Online]. Available: https://plotly.com/python/.

[39] I. Alexa Internet, *The top 500 sites on the web.* [Online]. Available: alexa.com/topsites.

[40] ISPQUICKLIST.COM, *List of internet service provider worldwide (isp) database*, 2021. [Online]. Available: https://www.ispquicklist.com.

[41] R. NCC, "Youtube hijacking: A ripe ncc ris case study," 2008. [Online]. Available: https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study.

[42] S. Janardhan, *More details about the october 4 outage*, Engineering at Meta, 2021.

[43] A. Mitseva, A. Panchenko, and T. Engel, "The state of affairs in bgp security: A survey of attacks and defenses," *Computer Communications*, vol. 124, Apr. 2018. DOI: 10.1016/j.comcom.2018.04.013.

[44] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the internet," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '07, Kyoto, Japan: Association for Computing Machinery, 2007, pp. 265–276, ISBN: 9781595937131. DOI: 10.1145/1282380.1282411. [Online]. Available: https://doi.org/10.1145/1282380.1282411.

[45] L. Salamatian, D. Kaafar, and K. Salamatian, "A geometric approach for real-time monitoring of dynamic large scale graphs: As-level graphs illustrated," *ArXiv*, vol. abs/1806.00676, 2018.

[46] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '99, Cambridge, Massachusetts, USA: Association for Computing Machinery, 1999, pp. 251–262, ISBN: 1581131356. DOI: 10.1145/316188.316229. [Online]. Available: https://doi.org/10.1145/316188.316229.

[47] A. S. Tanenbaum, *Computer networks*. Upper Saddle River, NJ: Prentice Hall, 1996.

A. Definition of Task

IAS | INSTITUTE FOR APPLICATION SECURITY

Visualisierung und Analyse von Knoten- und Kontrollpunkten im globalen Internet auf BGP-Ebene

Exposé zur Bachelorarbeit

Tim Christian Sauer

25. Januar 2022

Deutscher Titel	Visualisierung und Analyse von Knoten- und Kontrollpunkten im globalen Internet auf BGP-Ebene
Englischer Titel	Mapping and Analyzing BGP-level Hubs of Control in the Global Internet
Anmeldedatum	19. Oktober 2021
Abgabedatum	19. Januar 2021
Erstprüfer	Prof. Dr. Martin Johns
Zweitprüfer	Prof. Dr. N/A
Matrikelnummer	4971731
E-Mail	tim.sauer@tu-bs.de
Studiengang	Informatik

1 Introduction and Motivation

The Internet is a decentralized system built of a rapidly increasing number of routers. Several routers that share the same IP-prefix[1] and are under the administration of a common organization (e.g. Internet Service Providers, Scientific Institutions) form an *Autonomous System* (AS, Plural: ASes). ASes are identified by a globally unique number called *Autonomous System Number* (ASN).
For pathfinding between ASes, the *Border Gateway Protocol* (BGP) is used. To connect with other ASes an AS A announces its IP-prefix and ASN to its physically connected neighbors. A neighbor receives such an *BGP-Announcement* and decides if it prefers this way to exchange information with A. If it does so, it adds its own IP-prefix and ASN to the BGP-Announcement and passes it on to its own neighbors. By doing the process repeatedly, a BGP-Announcement grows in its number of ASes and becomes an *AS-Path* which represents a route between ASes. Note that an AS can learn different AS-paths from its various neighbors that lead to the same IP-prefix, at the end of the path. There are several criteria an AS considers when deciding which path it will use further. Usually this is the path length. This means that the route with the lower number of ASes along the path to the destined IP-prefix is chosen.
There already exists a large body of research regarding BGP and its security [3–6]. As BGP was first designed in 1990 there have been little efforts towards security at this time. Over the years there have been several incidents of Hijack and even Interception of Traffic, based on exploitation of flaws in BGP.
Let us assume a country whose government aims to spy on its citizens traffic. To do so it would have to attack those ASes which are used for routing the citizens traffic.
One could question if such attacks are even necessary if those targeted ASes are already under the control of the given countries government. This brings up our research question: Let there be a set A_{orig} of origin ASes and a set B_{dest} of destination ASes. AB is the set of all existing paths between all origins from A_{orig} and all destinations from B_{dest}.

Are there certain ASes which are along a high number of AS-Paths from AB and therefore state interesting points for observation of Internet Traffic?

[1] An IP-prefix is the part of an IP-address which identifies the network, a target device is located in. In other sources it might also be called *host-portion*.

2 Contributions

To approach the research question this thesis will give the following contributions:

1. Document the state of knowledge regarding the analysis of AS connectivity and important local and global hubs.
2. Implement a tool to visualize all ASes from given input sets A_{orig} and B_{dest} and the resulting AS-Paths from AB. Subsequently highlight those ASes which are intensively used along AS-Paths from AB and therefore state potentially sensitive points for observation.
3. Align the illustrated ASes by their geographic location to monitor which ASes are important hubs for local (inside a country) and which ones for global connectivity (between different countries).
4. Investigate if there are *essential* ASes, from those hubs discovered in point 3, which cannot be circumvented when routing traffic.

Although point 2 and 3 of our contributions have already been studied by Acharya et al. [2] we are curious to ask if their findings can be reproduced four years later. Over this time the number of ASes in the Internet has almost doubled [1, 2] and this growth might have also brought changes in topology.

Literatur

[1] World - autonomous system number statistics.

[2] Hrishikesh B. Acharya, Sambuddho Chakravarty, and Devashish Gosain. Few throats to choke: On the current structure of the internet. *CoRR*, abs/1806.07038, 2018.

[3] Hitesh Ballani, Paul Francis, and Xinyang Zhang. A study of prefix hijacking and interception in the internet. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '07, page 265–276, New York, NY, USA, 2007. Association for Computing Machinery.

[4] Asya Mitseva, Andriy Panchenko, and Thomas Engel. The state of affairs in bgp security: A survey of attacks and defenses. *Computer Communications*, 124, 04 2018.

[5] H. Roberts, D. Larochelle, Robert Faris, and J. Palfrey. Mapping local internet control. 2011.

[6] Matthias Wählisch, Thomas C. Schmidt, Markus de Brün, and Thomas Häberlen. Exposing a nation-centric view on the german internet – a change in perspective on as-level. In Nina Taft and Fabio Ricciato, editors, *Passive and Active Measurement*, pages 200–210, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

3 Implementation

This document is the foundation for the implementation of the *Bachelorarbeit* and describes the contents and goals that must be fulfilled. Please consider the general *Bachelorarbeit* requirements of your examination regulations. If you are in doubt please contact the examination office.

Task Definition and Supervision

Prof. Dr. Martin Johns

(Date, Signature)

Student

Tim Christian Sauer

25.01.2022, T. C. Sauer

(Date, Signature)